BUENOS AIRES – How It Works: Internet Standards Setting
Sunday, June 21, 2015 – 09:30 to 11:30
ICANN – Buenos Aires, Argentina

UNIDENTIFIED MALE: This is the inaugural instantiation of these tutorials, and the goal here is to help community members understand more about how the Internet fundamentally works. We're fortunate today to have Russ Housley who is a former chair of the IAB, has a consulting company, Vigil Sec, that does interesting work primarily in the security space as I understand it. He will be providing an introduction to the IETF and the assorted bodies associated with it.

With that, I will hand it over to Russ. Thank you.

RUSS HOUSLEY: Hi. I would like to keep this pretty informal, so if you have questions as we're going along, please feel free to ask them. For the recording, I will repeat them so the question and the answer get recorded.

David approached me and asked me to do this. Last March, I stepped down as the chair of the Internet Architecture Board, and I had that position for two years. Before that, I was the IETF chair for six years, so I have been in and around this process for quite a while.

I thought the first thing to do would be to explain where the IETF fits into the Internet ecosystem. This is one way of looking at the ecosystem. Not so much that anyone has tried to get consensus on a chart that describes the ecosystem, but given this one, which is

available on the Internet Society website, the part we're going to be talking about today is this part where the Internet Society, IETF, IAB, and Internet Research Taskforce are in this. Some people call it the Trivial Pursuit chart. Those of you who have ever played the game will understand that.

What is the IETF? At IETF 50, we had a T-shirt, and on the back of it, it said, "We make the net work," and that really is the one-sentence summary. But the mission is actually to produce documents, and those are documents that are intended to be high quality, technically relevant, timely that help people to use, manage, and do other things to make the Internet work better.

These documents are published in a thing called an RFC. The name is historical. RFC stands for Request for Comment. Many of them today, we're not asking for comments, but the name has stuck around.

The person who wrote the first RFC is well known in the ICANN community. It was Steve Crocker who is currently the chairman of the ICANN Board. He has the honor of having written RFC 1.

So the RFCs, the first one was published in 1968, and you can see that we've been producing a lot of them over the years. In 1986 is when the IETF was formed, and you can see on the chart where ICANN was formed as well.

Today the vast bulk of the RFCs, over 90%, come from the IETF. Although, some come from the IAB, some come from the Internet Research Taskforce, and some come from individuals who want to

make a statement about things in the Internet. That is what we call individual submissions, and there's an editor that you contact if you're interested in writing an RFC and not participating in any of these groups to get that done.

Some of the RFCS that have been around and influenced the domain name system – the first two on this list were before the IETF was even created. That was the creation of the domain name system and the structure and delegation for domain names. Then the IAB did a document that basically calls for there to be one root to the DNS, as opposed to everyone having a different view of how that system works. Then there's the WHOIS system and the means of signing DNS entries called DNSSEC, internationalized domain names.

Then there's the weirds. Working group created this registration data access protocol, which is used by registries and registrars. Recently, we gathered the principles for operating IANA registries and put them out in a document. The IAB did that.

Many of you may be aware that we're working on improving a bunch of protocols including: HTTP version 2, and then there's the PRECIS framework, which is about how to compare two strings that contain internationalized character strings. These are the kinds of things that are going on in the IETF that affect the community here at ICANN.

One of the ways to think about the IETF is from a cultural perspective. Dave Clark from MIT described it as, "We reject kings, presidents, and voting. We believe in rough consensus and running code." I think you'll understand this motto deeply by the end of the presentation.

But what it really says is voting isn't the thing that matters. Having agreement in a particular direction and working implementations that prove that direction gets the job done is the argument that will win the day.

This little blurb on the chart – that's a quote from RFC 4677 – tries to describe the IETF. It's just two sentences. That's a tough task.

But basically, people who come to the IETF are people who care about working on the technical side of making the Internet work better. Basically, everyone is invited to participate, and it's an open and international community. People who come are asked to participate as individuals not as representatives of their company. As I said, the goal is rough consensus, and we don't do voting.

One of the things that is important to understand about rough consensus is that when you're working on a technical problem, you have seen arguments where two people start saying, "I believe A," and somebody says, "No. I believe B."

You just watch the discussion of:

"It's A."

"It's B."

"It's A."

"It's B."

"It's A."

"It's B."

You're not making any progress. What will usually break that log jam is for one of them to say, "And here's a piece of code that proves my point," and once you do that, the other side either has to put up a similar piece of code to show that their perspective is right and that they were both alternatives, or the discussion is over.

Since the IETF has no members, it's actually impossible to have voting. But we have participants, and we also do not make any attempt to achieve unanimous agreement. That's part of the rough consensus part. It's okay if there's a few people who think another way would have been better as long as the vast bulk of people agree that a particular approach is good enough.

So we often, when we're having a debate, try and figure out what the sense of the room is, and we do this by humming. Humming is a way where you can very quickly learn whether the room is all in a particular perspective or deeply divided. A working group chair might say something like, "If you think that the direction that we've been discussing for the past ten minutes is the way that we should go, hum now for yes. Hum now for no. Hum now for I need more information."

Then very quickly the working group chair will be able to say whether the vast majority of people in the room fall into one of those three buckets or it's evenly divided kind of a thing. So that's the sense of humming that goes on in the IETF, and it really does very quickly work and without even taking the time to count hands.

In most cases, disputes are resolved through discussion. As I said earlier, sometimes they're resolved by posting a piece of code. Always after a discussion like that that leads to a hum, there will be something on that working groups mail list that says, "At the meeting, the sense of the room was that this is the way that we want to go. If anyone has concerns with that, speak now."

What happens there is it makes sure that people who were not able to attend the meeting are not left out. It also makes sure that you didn't get into a situation where you're having a group think going on, and after everybody left the room, there wasn't a realization that, "Oh. We didn't consider this aspect of the problem." So it makes sure those two things are taken into account.

There's an interesting RFC 7282, which is about consensus and humming. It makes the point that one of the things in achieving rough consensus is not a person with a strong technical view just being completely ignored. For example, you have not achieved what we consider to be rough consensus if there is a person who actually can show that, say, the thing doesn't work.

For example, if they can write a compliant implementation for one side and show that some part of the interaction moves to a corner case where the thing breaks. If that was to be the case, then you have not received consensus.

Anyway, it's a very interesting document. It's written for anybody. It's not full of IETF jargon, or the jargon that is in there is explained. So if

you're interested in how the IETF rough consensus works, then I think you would find that an interesting read.

Open standards is the real point of the IETF. However, when you think about the Internet as a whole, there is nobody who gets to say, "This RFC was finished. Implement it." No one is in charge. No one could dictate that.

Instead, the Internet works because a bunch of people are choosing to cooperate. What happens is when you have a gathering, such as is happening here in Buenos Aires for this meeting, each person comes to that meeting bringing a unique perspective of the Internet, and that diversity actually makes reaching consensus difficult. That's why there are the discussions.

Yet, at the end of that discussion, when consensus actually does emerge, the outcome is better, having considered all those perspectives. It's also clearer because people have had to take the time to write it down as clearly as possible to get the people to agreement, and it's more strongly supported because you've changed the minds of some people or worked through issues that the original proposal didn't even know were there. So that process has created a better document, support for it, and a better description, and that's very valuable.

It also means that when somebody didn't participate in that discussion, the better document is there for them to decide whether they want to implement it or not.

An IETF document: they're always in English. One of the things that the IETF did – actually it was two chairs before me, so fortunately, that person was a Norwegian – led a discussion about what language the IETF should do their work in. There was a very long and heated discussion, and the answer was we're going to do it in one language and one language only, and that's going to be English.

However, all of the documents that are produced by the IETF include a blanket license for translation to any language. So while we're not doing the translation ourselves, anyone who wants to do so already has permission from the beginning.

In addition, the IETF uses ASCII for its documents right now. Internet drafts, which are things like the working documents that are targeted to someday possibly become an RFC, are ASCII, and RFCs are ASCII.

We're in the middle of a long transition to move away from that to XML. XML will allow us to have a more rich environment to include figures – other than little stick figures – and from that XML, we'll continue to produce ASCII documents or HTML or PDF or any other thing that emerges in the future. People are already beginning to talk about e-book formats for example.

So, this is going to be a long transition because it's important to us that you continue to be able to read the documents many, many years after they're written.

For example, RFC 20. It was written by Vint Cerf. It says, "The Internet protocol should use ASCII." But you can still read it today, and that's

the important thing. We want that to be true of the documents that we produce in this new format. Will they be able to be read by anybody 40 years later?

I was trying to describe the Ethos of the IETF on one slide. It's actually not as easy to come up with it as when I started. But I did this a couple of years ago when I was asked to explain the IETF in under five minutes to a group. Anyway, so this, I think, is as close to a summary as you can get.

The first is that it's open standards and open process. All people who have any interest in the topic at hand are invited to participate. Even if they cannot make the face-to-face meetings, they can join the discussion on the mail list, and all documents, whether they be the final specification or any draft leading to it, are online, available for free to anyone.

That sets the who's there, and we're not going to hide any information. It's all open and transparent.

But the goal is one Internet. That is, we're trying to come up with standards for a global Internet, not little pockets, which means we're going to maximize the interoperability and maximize the scalability, and we're going to avoid specialization in different places, whether places be geography or any other dimension of how you might subdivide the Internet. We, the IETF, are trying to work on the whole.

Contributions are sometimes judged harshly, but they're judged on the merits. As I said, rough consensus and running code. You'll often

see at an IETF meeting, some folks having a pretty rigorous argument about some technical detail in a proposal, and when all is said and done, the meeting breaks up, perhaps neither one won or lost that day, but they'll go have a drink in the bar. It's about getting the technology right, not about the human interactions there.

This graph shows the number of people who have been coming to IETF meetings since the beginning. The very first meeting, IETF 1, there were 21 people there – small group. The Internet was mostly universities connected to a research network, and they were working on some of the very same problems that we continue to work on today. Some of them have not been solved yet. How do you scale routing to global constituents? How do you make sure you never have collapse due to congestion issues? Those kinds of things.

At the peak, we were just under 3,000 people at one meeting. What happened that caused us to fall off that peak was the .com bubble burst. Like many of the organizations in the Trivial Pursuit chart I showed in the beginning, that economic downturn had an effect. At the same time, it never fell below 1,000 people, so it's not like it totally became uninteresting, which that's still quite a lot of people working on protocols to make the Internet work.

I thought I'd talk for a minute about when does the IETF decide to do a particular project, or basically, when is the IETF the right group to do the work as opposed to someone else on that chart?

The first thing is that there's a problem that needs to be solved, and again, it's going to avoid the specialization of protocols in different

parts of the Internet or different parts of the world, and is the research complete? That is, is this a research problem or an engineering problem?

If it's a research problem, it probably goes to the IRTF. But if the research is complete and only the engineering is left to be done, then it's the IETF.

Is the scope well defined? This actually was a problem when some people came to the IETF and said, "Well, let's work on cloud."

We said, "Okay. What part of the cloud problem are you trying to solve?"

Cloud is too ill defined to just have a cloud working group. You want to know what the deliverables are going to be, and by asking that question, you end up scoping it. We did do work that supports cloud, but we didn't do a cloud working group.

We also want to make sure that there's enough people who are willing to work on the project and that those people believe the scope of the project is such that it can be completed in a timely manner. We don't want to work on things that are open ended and go on forever and you don't know when you're done. That's just the nature of the IETF working group.

The IETF is most successful when the participants that are engaged really care about solving the problem, and it works best when people come with an open mind and consider the other peoples points of view. I suppose that any of these collaborative processes work best

when people are interchanging ideas openly and without personal attacks. I guess that's what I'm saying here.

However, the IETF has had a couple bad experiences where we have collaborated with other standards development organizations. It turns out, pick any two processes, and if you try and then get approval for two documents using both of the processes and neither process terminates until you come to exactly the same language in both documents, you're just setting yourself up for deadlock.

What we've done more recently is taken a different approach. When collaboration is needed, we figure out which part the IETF is going to work on, which part the other organization is going to work on, and have a clean interface between the two.

An example of this is the W3C and the IETF collaborating. One of the specifications is called WEBrtc, and the other specification is called RTCweb. What they are is one is a protocol specification, and the other is the APIs in the browser environment that are necessary to implement the specification.

There's an example where in the collaboration, the work is cleanly divided so you don't need both the W3C process and the IETF process to agree on exactly the same words at the same time. The system isn't complete until both sets of work are done. We find that kind of work to be far more successful and certainly much less onerous on the participants.

Like any organization, the IETF has a structure. One of the interesting things is the IETF is not a legal entity in its own right. Instead, we are an activity within the Internet Society. That was a decision that we made several years ago. The IETF actually existed before the Internet Society and so we had to make a decision to become our own legal entity or find a parent, and we chose to find a parent.

Right now, there are about 1,000 to 1,200 people coming to each of the three meetings a year. There are many, many, many more people participating on the mail lists. Last time I looked, there were about 60,000 unique e-mail addresses on IETF mail lists. That's a huge number, and you can't readily tell whether one person has multiple e-mail addresses and has one associated with each working group because of the way they do foldering or something, so we don't really know exactly how many people participate in the IETF for that reason.

Right now, there are about 120 working groups. This is the place where the real work of the IETF gets done. We are constantly creating new working groups to focus on particular activities and closing working groups when that piece of work is finished.

The working groups are divided into six areas. Each of those areas is led by two or three area directors. There is one exception, which is the general area, which is led by the IETF chair. The general area is the one area that does no technical work. What the general area does is manage the IETF processes.

We use the same working group model to manage the IETF process as we do to do the technical work. There's just not as much going on

there. In fact, right now, there is one working group in the general area, and I have several slides about it at the end because that working group is working on the IANA stewardship transition.

Other management bodies are the IESG, which is the Internet Engineering Steering Group and that is made up of the area directors. That is the group that does the final approval for any document that's going to become an Internet standard.

The IAB provides architectural guidance. It is guidance. It is not a dictate by any stretch. We manage liaison relationships between the IETF and other standards organizations, such as the W3C, the IEEE, ITU, the International Standards Organization (ISO). We have some relationships with Etsy and so on.

Then there's the IAOC, which is the IETF Administrative Oversight Committee. It's where the budget gets done, contractors are set up, for example, for our secretariat, and where the intellectual property rules are dealt with.

This whole structure has one employee, and that person technically is employed by the Internet Society. But everything else is handled by contract, and that one employee's job is to manage the contracts.

So this is a way to look at it. It's kind of difficult to say. It's not a hierarchy. That's the first thing.

So the hierarchy is the IESG – in the blue boxes – each having their area and the working groups under them. The rest are a supporting

constellation – the Internet Society providing the organizational home and being the legal entity where the checking account can live.

The IAB – actually in addition to providing the things I said on the previous slide – provides oversight for important things, such as IANA, at least the protocol parameter part of IANA. It provides oversight for the RFC editor, and it is the final body in the technical appeals chain.

If someone has a concern with a decision that is made, the IAB will be the last body to have a say on a technical decision. In addition to that, if the appeal is not about a technical question, but is about a process question, basically, the person raising the appeal is saying, "You didn't follow your own documented processes," then there is one step beyond the IAB which is the Board of the Internet Society.

Other things on this chart: the bubble that says IASA in it. IASA is the group that manages the budgets and that kind of thing, and the two subparts are the oversight committee and the IAD is the IETF Administrative Director. That's that one employee.

The IETF Trust was set up at roughly the same time that the IETF was making the transfer into this structure. The trust is a legal entity, and it is the thing that holds the copyright for the RFCs, and it owns the domain names that are used in the structure – IETF.org, IAB.org, IRTF.org, RFCEditor.org. Those – prior to the creation of the trust – were held by CNRI, which is an organization that was contracted way back prior to IETF 1 to manage the IETF secretariat.

When the U.S. government stopped funding that, this is when the new structure was set up. So, it has to do with the transition of the Internet from the days of being a research network funded by the U.S. federal government, to the open commercial Internet that we all know today. Next.

I think I said all of this when talking about the last chart, but the Internet Society, or ISOC, is the administrative home for both the IETF and the IAB. The Internet Research Taskforce is actually under the IAB. That's one thing that I didn't know how to show on that chart.

As I said earlier, the Internet Society was formed when the National Science Foundation in the U.S. decided to stop funding the IETF secretariat. One of the other things that the Internet Society president does is – at the beginning of the IETF process – to select people to hold leadership positions. That whole process is started by the president of ISOC choosing the chair of the IETF Nominating Committee. That person has no vote in the process, but that's how they are picked.

The IAB provides advice about the global Internet architecture. It also provides technical advice to the Internet Society. It turns out that this advice usually requires the IAB to meet for several meetings or teleconferences to provide. If the question is small, it usually doesn't come to the IAB. When the Internet Society wants technical advice from the IAB, it's usually a more deeply considered question that has many different pros and cons to the alternatives.

As I said, the IAB manages external relationships for the IETF, and as part of its oversight role over the RFC editor, it actually names the

person who is the RFC series editor. The IAB selects the persons to run the IANA registry for the protocol parameters, and once that is selected, oversees the operation. Presently, that is the IANA department within ICANN, and it has been that way since the creation of ICANN.

The IAB is part of the selection of the leadership. When the nominating committee comes up with the membership for the IESG, the IAB acts as the confirming body for that.

As I said, it's the top of the technical appeals chain, and there's one step beyond that for the ISOC Board if the appeal is about process.

These six areas are the areas within the IETF that make up the steering group, and all of these have working groups. As I said before, the general area is led by the IETF chair, and the other five have either two or three area directors. The ones that have three just simply have a little more work going on in them, and so in order to do the steering for those, they just have spread the load among three people instead of two.

The area directors are responsible for guiding the work within the area. Not for doing the work themselves by any stretch, but they are responsible for managing the process. For example, they appoint the working group chairs.

Most working groups have two chairs. A few have one, a few have three, but most have two. The area directors collectively as the

steering group, charter the working group, but the area director who is responsible for that group once chartered will name the chairs.

The area director can close a working group whenever that person feels the work is complete or if they feel that the work will never be complete because the group has lost focus.

The area director will review the documents that are produced by the working group for basically three things: to make sure that a sound technical solution has been developed. If they don't believe that, then they can't be an advocate to the rest of the steering group for the solution.

The second is, was proper process used, and the third is, is the thing that was actually produced what was called for by the charter of the working group? Did they actually do what they said they were going to do? Those are the three questions that an area director will ask when they get a document.

Those area directors – 15 of them total – make up the steering group. They are the ones who, when the working group is done with it, the area director for that working group will bring the document to them, and they are the final approval. Once they say, "Yes," it goes to the RFC editor for publication.

They also, as I said, are the group that does the working group creation. That is, these 15 people have to agree this is a piece of work the IETF should be doing.

They review the solution because you want to make sure that all aspects of the solution have been considered. So if the work got done in the security area, you can be pretty sure it's a sound technical solution, but does it impose anything on other parts of the Internet that you didn't anticipate? That's what this multi-discipline technical review is to catch, and sometimes some things are caught at this aspect.

It's not the best time to catch a problem because, as in any specification or development effort, the end is the wrong time to say, "Hey. There's a problem here. Let's go back and figure it out."

But at the same time, it's difficult to get people from other disciplines to do a document review if it's a very early draft because they know it's going to change. "Why should I spend my time on that yet?" So, there's a tension there, and this is the solution to make sure it has a look from people who are looking at the Internet from different perspectives.

In addition, documents that come from the research taskforce and from that individual submissions mechanism that I talked about earlier go here to have one question answered about them. Is this document really an attempt to "end run" the Internet standards process? If the answer to that is no, it goes on and that's the end of it. If the answer to that is yes, then there's a big discussion, and that big discussion usually leads to you need to bring this document to the IETF process.

Yes? Sure. The question was what do I mean by end run? Let's see if I can come up with a real example. Yes. There's a thing called TFTP, which stands for Trivial File Transfer Protocol. It is used by some low-end devices to download firmware, for example.

An implementer says, "I can make TFTP run faster by increasing its window size." What a window size is, it's the number of packets that can be sent from the server to the client without yet having an acknowledgement from the client. Does that make sense?

So, TFTP was really stupid. The window size was one. So send me a packet. Got it. Send me a packet. Got it. This person had suggested that the window size, if it was bigger, the protocol would run faster because the latency from the time the server sends it until the client gets it could be consumed and it would run faster.

The way that they did this actually caused a problem for congestion control because they didn't do it using TCP. They made up their own mechanism.

That's an example of you didn't consider congestion control. That's an IETF topic. We need to make sure, and so the author and the people in the transfer area had a long discussion about how do you make sure that the server doesn't send so many packets against that client that they cause a congestion problem. It turned out to be a long discussion because the person said, "It works fine on my LAN. What are you talking about?"

We said, “If you put a couple routers in between it, what’s going to happen?”

He says, “I don’t know. I never use TFTP in that environment.” You can imagine that’s where the discussion went. Did that help? Good.

So, working groups. As I said several times, this is the primary place where the real work gets done. The IESG is the approval for the creation of them. As part of that, the IAB is asked for perspectives. Sometimes we have them; sometimes we don’t.

The working groups tend to be short lived. By that, I mean three to five years. But they are chartered to work on a specific problem with specific deliverables. There’s no formal membership. Everyone who has an interest in that topic is invited to participate. Every working group has its own mail list to have that discussion.

The working group chair sets the agenda, appoints the editors for the document, and optionally appoints a working group secretary. The most important of their jobs, in a sense, is they’re deciding when they’re done. They are the rough consensus caller, not just for “does this working group believe in this document,” but lots of the little decisions that led up to that as well.

This is the way to look at the approval process that I’ve talked about. Working group actually or an individual who believes that a document should be a standard submits it to a sponsoring AD. I would say, the vast majority of documents come from working groups. Although,

sometimes the individual proposing something works fine, goes to the sponsoring AD, the sponsoring AD does an IETF-wide last call.

For documents that came from working groups, that's two weeks. For documents that came from an individual, that's four weeks. The reason for the longer time is the assumption is that since the working group didn't work on it, we ought to get more eyes on it.

Those last call comments come to the IESG as a whole. They are considered as well as their own review, and once it's approved, it goes to the RFC editor.

If those concerns are not small, it can go back to the working group and say, "These issues were found. I think you need to address them." In which case, it would start over.

One of the things that we will do when trying to decide if we want to have a working group at all to help identify the need is hold a Birds of a Feather session or a BOF. Basically, this is a time at an IETF meeting, usually two to two-and-a-half hours where people will put forward the idea for a working group, and then the community will have an opportunity to discuss it. BOFs tend to be very well attended.

The group is meeting to really answer three questions: to show or demonstrate that there's a need for the thing that's being proposed, to show that there are people interested in using the outcome, and to identify people who are willing to do the work. Because if you have the first two, but no one cares enough to work on it, it doesn't ever get done.

One of the results of a BOF is often a draft charter that can then go to the IESG for consideration. Most BOFs have a mail list as well, which if chartered, becomes the working group's mail list. The working group then will produce Internet drafts. As I said earlier, all drafts at any level are public and available on the IETF servers for download by anyone.

Very often, when you have a big group working on something and you're trying to resolve a small technical piece of the problem, the working group chair might have a design team to go off and bring back a recommendation to the group as a whole. When that happens, the group will discuss the design team result on the mail list as well as at the face-to-face meeting.

When you get to the end, most working group chairs do a working group last call just to make sure the people who have been doing the work have a last call. They know we're coming to an end. If you have anything else, raise it now.

Then it goes to the area director for review who initiates the IETF-wide last call. Then the IESG will evaluate and ultimately approve the document. When it becomes an RFC, the RFC editor posts it.

The IETF copyright is somewhat unique in the sense that the authors and the IETF both hold copyright. Many organizations that develop standards insist that the authors give their rights to the organization. For example, the IEEE does that. The IETF, in an attempt to make it easier for academics to participate, allow the authors to retain copyright.

One of the things that I always get asked when I do discussions like this is, "If there's no membership and no voting, how do you pick your leaders?" It's complicated, and it has to do with this nominating committee.

The chair is selected by the Internet Society president. What they do is they reach out to the community and say, "I need to get a pool of volunteers who are interested in being on this year's nominating committee."

The nominating committee is made up of ten of those volunteers. The pool of people who are eligible to be on the nominating committee is people who have attended three of the last five IETF meetings. We usually get between 100 and 120 people who are interested in serving on the nominating committee. The NomCom Chair selects ten of those at random in a verifiably random and confirmable way.

Here's how that works. We have an RFC that specifies a random algorithm that requires a seed. The seed is chosen by the NomCom Chair in advance of what the value will actually be.

For example, the closing of a particular stock price on a particular date, the size of the U.S. national debt, a lottery result. Several of these kinds of values are picked to become the seed, then anyone in the community can run the algorithm and see with the pool of volunteers the ten that are picked.

There's one complexity to the picking, not that that wasn't complex enough. That is that, at most, two of the people pulled from the pool

can be from the same organization. That's a mechanism we use to make sure that the NomCom is not captured by any of the particular organizations.

A couple years ago, we found that a third individual was selected where two previous ones were from the same organization, so you continue the algorithm, and you draw the 11th, and they replace that person.

The NomCom then has ten voting members, the chair, and the past chair. In addition to that, they get liaisons from each of the leadership bodies. Those liaisons are people whose terms are not up so that they have a way to ask questions about how the body actually does their work and those kinds of things. Then the NomCom will select a person for a two-year term.

The IETF chair, area director, IAB member, and two of the seats on the IAOC are selected in this fashion. Once the NomCom makes their selection, there's a confirmation step. The IETF chair and the area directors are confirmed by the IAB. IAB selections are confirmed by the Internet Society Board of Trustees, and the IAOC selection is confirmed by the IESG.

That's the way that the nominating process works, the selection process is confirmed and, ultimately, how leaders are picked in an environment where there are no members or voting.

Given that we're here at ICANN and that we're in the middle of this IANA transition process, I thought you might like to know a little bit

about the protocol parameter registries that are managed by IANA. These registries are full of values that are used by the Internet protocols. The tables of values are maintained by IANA based on policies that are set for each registry. Those policies are in RFCs, and so in a sense, the IETF is responsible for setting the policy for each of the protocol parameter registries.

ICANN has been doing that since 2000, and we have a Memorandum of Understanding that describes that relationship. That is captured in RFC 2860. Over time, this process has been evolving. We do this by updating a Service Level Agreement.

The IAOC and the IAB have a role in doing that. We have a joint group called the IPROC that is made up of people from the IAOC and the IAB who work with people from the IANA department to come up with that Service Level Agreement each year. Every two months, we look at the way the work has been done and say, "Is it meeting the Service Level Agreement?"

I have to say, ICANN is doing a stellar job. Their goal is to meet 95%, and I can't think of a time in the last couple of years that it has fallen below 98%. It's most often at 100%. Elise Gerich and Michelle Cotton do a great job.

What is a protocol parameter? IETF standards need things like port numbers or similar values, MIME types and so on. There are literally thousands of these registries.

This is not a unique task. Any group that does protocol standards has a need for these things. In the IEEE space, they have a thing called the RAC, the Registration Authority Committee, and they manage it for the IEEE. It turns out that they do it with employees of the IEEE, but that's how they do it.

Another example is an error code for the hypertext transfer protocol. That little gray box on the slide shows an example of some of the error codes and which RFC defined that error code and what it means. These are the kinds of things that happen.

Some of the registries have very few possible values. An example is for the IP protocol. There's a version number in it. It's four bits, which means it has at most 16 values, so we're not giving those away very easily. There's very tight controls on assignments of new ones.

Other things have an infinite possible number of things. An example of that would be an object identifier within the IANA arc. There are literally an infinite number of them because it's just an integer and it doesn't have to fit in any particular number of bytes. Those we pretty much do first come, first serve. Anybody who wants one can have one. There's two extremes.

There is no direct impact by the assignment of one of these values, and the reason is these values are used to implement the protocol. So the programmer who is doing the implementation goes to these tables, looks up the values, and uses them in doing their implementation.

It's quite different than the DNS root zone, which is published and the clients are making direct references to the database and the IP address that a particular person sees at the particular time is dependent on what's in that database. This is quite different. Because the programmer is in the loop, there's no, "I make this one change and every computer on the Internet has a direct change." That doesn't happen here.

How do we divide the labor up between IANA and the IETF? The IETF makes the policy decision. In fact, it asks that a particular registry be created, provides the initial set of values, and says, "Modifications to these values or additions of new ones have these policies." These are the steps that one has to go through.

The IAB provides oversight, and the IANA department within ICANN does the implementation by actually making the changes to those tables, publishing them at IANA.org. This arrangement has matured over and over. This basic structure of the three parties goes back to the beginning of the IETF, and we have a lot of good experience. And, as I said, with the updating of the SLA each year, the system just gets better and better each year.

Now we have this stewardship transition going on. I'm sure that everybody is aware of the contract between ICANN and NTIA. Then March 2014, NTIA says, "I think that we're going to step away from this role," and requested a transition proposal be provided by the Internet community for how that would happen.

From the IETF's perspective, this is what was done. Obviously, this hasn't completed, and this week, it will have a fundamental say on what actually gets delivered on the NTIA box on this chart. But we have had RFCs that define all these MoUs or the policies for the registries.

We have the MoU. We have the Service Level Agreement that's updated each year. These were inputs to the IANAPLAN Working Group.

The IANAPLAN is this active working group in the general area, and they have one task: produce the proposal for the ICG. The ICG being the coordinating group for the IANA transition. It produced the plan. There was an IETF-wide last call. That went to the IESG. The IESG has approved the plan.

It has not been published as an RFC. At least not yet, because when the plan from the numbering community and the names community make it to the ICG, it's possible that that analysis will find a conflict among the proposals or a gap, and if that requires a change in the proposal from the IETF to resolve that, then that will go back to the IANAPLAN Working Group and we'll do this again. Hopefully, that won't be the case, but that's the process that we've been following.

In fact, the approach that the IANAPLAN Working Group took was: what is the least amount of change that is necessary to the existing structure we've been using to address the implication that NTIA is moving out of its role? That is basically the question that drove the

discussions in that working group. Basically, minimal change to the oversight was the preferred way forward.

The working group requested that the IAB and the IAOC produce any necessary documentation to implement this agreement. They found the pieces of the NTIA contract that they hoped would be captured somewhere, and they asked that two changes be made to the SLA going forward that would cause that to be part of the agreement between ICANN and the IETF, and that's what's in the proposal.

At the bottom of this slide, you'll see a URL. That is the same place that all Internet drafts are stored. You can see, it took us ten iterations to get to the document because the first document is numbered 00, and the one that was sent to the ICG last January was numbered 09.

There was a realization that the U.S. government has very little role with protocol parameter registries, and so the IETF was able to put forward a very straightforward proposal. In fact, it was just pretty evolutionary as opposed to revolutionary. It just asked for these two provisions that are currently handled in the contract to be put in the SLA.

I hope that helped understand how the IETF is approaching this transition. I want to move to a bit of a summary.

Today the Internet works on standards. The Internet wouldn't work if clients and servers were speaking different protocols or even different versions of the protocol. The IETF uses a very open and inclusive and transparent standards process. Everyone's invited. All the documents

are there for anyone for free, and anyone with an e-mail address can participate.

Again, we're looking for one Internet, a global Internet with maximum interoperability, maximum scalability, and we want to avoid specialized protocols in different places.

I want to show a very brief movie clip – it's about a minute and a half – that I think did a really good job of taking all of that stuff and summarizing what the IETF is about. Let me play that for you.

[video begins]

The Internet is like a global city. We build it and use it together. But a vibrant city needs well-engineered infrastructure to keep it growing and running smoothly.

That's why we need the Internet Engineering Taskforce – a global organization of volunteers collaborating to design standards that provide the infrastructure for innovation on the Internet. The IETF is open – open participation, open processes, and open standards.

Today, because of the IETF, we can do so many things that we take for granted: e-mailing colleagues, instant messaging friends, and making phone calls to family, all over the Internet. E-mail and Internet addresses can use many different languages, and we can be more confident in the information we access over the Internet.

Tomorrow, because of the IETF, billions more people will have an address on the Internet, joining the two billion of us already

connected, and that means more people communicating and collaborating on new ideas. New ideas like an Internet of things connecting devices in your home and beyond over a smart grid that links everything from your thermostat to medical devices. Imagine all of the things the Internet will connect: your car, your appliances, and your mobile devices.

The IETF is open to you and your ideas. Support the Internet. Support the Internet Engineering Taskforce.

[video ends]

RUSS HOUSLEY: My understanding is that these slides are going to be made available on the ICANN website. They have my contact information if you need any follow up.

I want to point out a couple things. On this table over here, there are two handouts.

The first one is a brief thing about the Internet Engineering Taskforce in English and in Spanish, and the second is the announcement about IETF 95, which will happen here in Buenos Aires in April of 2016. That will be the first time that the IETF has a meeting in South America.

Our next meetings are in July in Prague, Czech Republic. After that, we're meeting in Yokohama, Japan. The handout is here about our first time in South America here in Buenos Aires.

The IAOC is responsible for picking the venues for where the meetings will be because it has tremendous budget implications. That's why the people who do the budget make that decision.

The Internet Society has for the past couple years been running a fellows program, which brings people who do government regulation and don't otherwise attend IETF to give them an exposure to what goes on at these meetings. It's been a very successful program and has brought people from around the world to see the IETF up close and personal, witness a hum, and actually see that the people that come to these meetings are deeply technical and able to participate.

There's a description of that fellows program on the Internet Society website, and I know that they have a lot of people who reach out to them and try to come. If you know a regulator who might really benefit from seeing how standards are made, you might want to point that to them.

Any questions whatsoever? Yes?

UNIDENTIFIED MALE: One of the things that you mentioned in your presentation is the Internet Research Taskforce. This is the first time I have heard of something like that. Can you elaborate on how an organization or one can join the IRTF?

RUSS HOUSLEY: Yes. The question, in short, was how can you get involved in the Internet Research Taskforce?

The answer is it depends. I'm sorry that it's that complicated, but different research taskforces are set up in different ways.

Some of them are completely open and totally available to anyone who just wants to join the mail list and get involved. An example of that is the CFRG, the Crypto Forum Research Group. If you're interested in cryptographic algorithms that might help solve problems on the Internet, that's a group for you.

Other groups are very closed. A group of researchers are working on a tough problem, they want to make sure that they have consistent voices as they work through a particular thread, and so they reach out, get what they think is critical mass, close the group and work together.

An example of a group that did that was the folks that did the fault-tolerant work that is used for sending IP between spacecraft and Earth. They felt that if they had an evolving membership, they would just never get done or certainly not before anything got launched that they wanted to be on.

Everything in between happens, but that's the extremes. If you go to IRTF.org, you will see a list of all of the research groups that are currently going on.

The IRTF chair is appointed by the IAB. That's currently Lars Eggert from NetApp. He has tremendous discretion. If a bunch of people

come to him and show that there's a problem and people willing to work on it, he basically will create it.

Like the IETF, they have a steering group. The difference is that steering group is made up mostly of the research group chairs and a few IAB members. Does that help?

JOSEPH MARC: Good morning. Thank you for the presentation. It's very useful.

My name is Joseph Marc. I'm a returning fellow from the ICANN.

Actually, I have two questions. The first one is about the NomCom, the Nominating Committee. What is the background, the requirement for the people which will join the NomCom?

RUSS HOUSLEY: Let's do them one at a time. That's very simple, straightforward. If you want to be a voting member of NomCom, you have to have been to three of the last five IETF meetings.

To be picked to be chair of the NomCom, basically, people who have achieved that have been long-time participants. The current NomCom chair is Harald Alvestrand who was IETF chair three ago. Once you're part of the leadership, you never get away.

JOSEPH MARC: The second question is about DNSSEC. I would like to know what the input that IETF has brought in the contribution of the DNSSEC [and direction]. Thank you.

RUSS HOUSLEY: The specifications for DNSSEC were done in the IETF. The two RFCs that I put on that early chart – I can back up to that – those two RFCs (4034 and 4035) specify how DNSSEC works. It specifies what bits get stored in the DNS, and there's an IANA registry associated with that that says which cryptographic mechanisms are used for signatures that include the hash function of the digital signature algorithm. One of those two establishes that registry, and we know that a couple other entries have been added for how, if that algorithm were to be deployed, this would be the code point for it and so on.

After those specifications were done, there was a slow, long process to figure out how they would actually get deployed. Some of the early adopters like were.se, where they started signing their zone, even though the root was not yet signed. Then ultimately, the root zone was signed, and now many, many, many people are being encouraged to, now that the top is signed, let's see if we can get the whole hierarchy signed. Does that help?

Any other questions?

UNIDENTIFIED FEMALE: Does the IETF do any tracking of the implementation or application of its RFCs?

RUSS HOUSLEY: Yes and no. Nothing formally.

I would say the people who work so hard to get through this process I've described pretty much do know whether their work got legs. We wanted to see whether we could actually measure that in any credible way, and the first step was to create a Wiki where people could tell us about their implementation that does a particular thing. So right now, what we have is an extremely voluntary way to find that out.

There have been other efforts to pay attention to which protocols are being used by looking at port numbers and other things like that through various observatory programs.

The advent of tunnels and VPNs hides a lot of that, and so we don't really have good data from those efforts. Other big applications like browsers sometimes include telemetry that reports back information – some of which is made public, some of which is not – and so through that, we have some insights. But do we have an absolute understanding? No.

UNIDENTIFIED MALE: When writing a draft, how do you go about selecting a working group? I'm doing an RFC at the moment. I'm struggling to find the correct working group to categorize it in.

RUSS HOUSLEY: Ah, I see. It really depends on the topic. Maybe we can meet later and I'll offer you some advice. But sometimes there is no working group. In which case, you need to do the BOF process to determine whether there are others who are willing to go on the journey with you.

Other times there's a working group that's working on the problem that your idea is perfect for, but they may be way late in the process. If they're about to do working group last call, the last thing they want to hear is, "I've got a different idea." You may not be welcome in that kind of an environment.

Sometimes you don't want a working group at all. Sometimes the problem is very narrow and the number of people who are going to pay attention to it is small enough and you know them anyway. What you want to do is go to the area director and say, "Here's my idea. I've had these ten people look at it. You know them all. You can tell they're interested, and maybe we ought to go straight to last call."

There's a very diverse answer there, but I'd be glad to listen. I don't promise that I will know the answer, but I can point you to someone if I don't.

UNIDENTIFIED MALE: Any last questions? All right.

Well, thank you very much, Russ. Thank you to all the attendees for being here at the first of these tutorials. We hope to continue this on. Make it better, stronger, faster – or maybe not faster. At subsequent ICANN meetings, you can assume there will be a similar set of

tutorials, probably a bit expanded in the Dublin meeting, so I would encourage you to let your friends and neighbors know.

If you're interested in these sorts of topics, we have a couple more tutorials today. One on basically how the Internet protocols themselves actually work and then a more specialized tutorial on how DNS registry protocols work. Those are later on today in the same room.

Thank you for attending. I guess we have about a half an hour or so until the next tutorial. No. Lunch. Time flies. Thank you very much.

[END OF TRANSCRIPTION]