BUENOS AIRES – RSSAC Workshop
Wednesday, June 24, 2015 – 15:30 to 16:45
ICANN – Buenos Aires, Argentina

UNIDENTIFIED FEMALE: …from the IAB, and we have a liaison from SSAC and we have a liaison to the board; Suzanne, who just introduced herself. And also a liaison to the Nominating Committee that nominates board members.

Before we continue, I'd like anyone from RSSAC and the caucus to stand up just to be recognized please, and the caucus please.

I'd like to move on to briefly talk about the caucus and its purpose. As we said, we are a tightly scoped body, a very small group with a very defined function. We are an advisory to the board. They come to us and ask for our opinion on technical matters. And we call upon the caucus to do pretty much a very large chunk of the work.

And the RSSAC is contained in the caucus, so if you look at this from a set perspective, the superset is the caucus and the RSSAC is a subset of the superset.

We call upon this body of experts to do our work, and as I said, it's a broad spectrum of expertise from security to DNS, and many are actually root server operators as well.

The body of work culminates in a document that gets published. We have document leaders. The document is scoped. They have deadlines. There is absolute transparency in what we do.

The caucus currently comprises of 61 technical experts, and we had a kickoff meeting in Dallas at IETF 93, and 43% of these caucus members are not from the root server operators community. They each submit a statement of interest, which is publicly available. We also credit their work as they make their technical contributions.

As I said earlier, we did kick off the caucus at the IETF meeting this past in Dallas. 40 caucus members participated and we talked on ongoing work, brainstormed on future work, and there is a process to join the caucus. Please send an e-mail rssac-membership@icann.org. What you need to do is send in an expression of interest and there's a committee that reviews it and moves the names forward to the RSSAC, and eventually there's an appointment process. I'm sorry it's not very clear. The slides are a little muddled.

We do publish documents in the form of reports and statements. For example, I'll talk about the first document there, RSSAC 001. These are service expectations of the root service. These documents are available on our website, and the goal is for – this is for public consumption and for the root operators to review these documents and see how they can meet those expectations.

We also submit statements that are published on our website. Recently, the CCWG presented some of their work and they had a public comment process, and we submitted our statement which is now available on our website as well. Suzanne will talk about that a little bit later.

I'll turn it over now to Jim Martin who's going to talk about RSSAC 002, one of the works that the caucus completed for us.

JIM MARTIN: Thank you. Good afternoon. Next slide. So RSSAC 002 is a document that was published late last year, was the work of some previous entities within RSSAC, and more recently a finalization by the RSSAC caucus. It was a body of work that was all based around the idea of wanting to understand the detail of how changes in the root server system will impact the actual servers. To that end, there were a number of metrics that we were attempting to collect.

So this document identifies what metrics, what statistics, should be collected from all the root server operators so that we can understand the change of the root server system over time. The idea is to use it as an early warning system, so if we begin to detect changes in the way that the root system is being used, that we can take action before it becomes a critical situation.

The measurements that are being requested in this document are how quickly the publishing data actually makes it from VeriSign generating the data all the way through to the furthest publishing server. The size of the zone, the number of queries that are being received per second, and this is [TCP, UDP], what's in IPv6, what's in IPv4, all these sorts of details. The query types and the response size distribution, which request codes are coming in.

And finally, as sort of a stretch goal is the number of sources seen. Oddly enough, that's some of the harder ones to do. If you go to the next slide.

What we've made the recommendation is that the RSSAC has requested each of the root server operators. We recommended to each of the root server operators that they implement the measurements that are requested in this document, and that those be published in such a way that the public in general has access to them.

We then are monitoring the progress of how each of the root servers, whether they're actually publishing and whether what they're publishing is correct. Finally, there is an ongoing process for reviewing these sets of statistics to see if they're still useful. Are we collecting data for a good reason? And is there something additional we should be collecting? That's an ongoing process every two years that cycle takes place. Next slide.

As of this morning, we went through and got an update on almost all of these letters and this is the current status of the data collection for RSSAC 02 data. So a number of the letters A, H, J, K, L all are currently collecting and that information is available. I'll show you where in just a second.

The rest of us are all in various phases of deployment, but the intention is that, from the root servers, they've informed us that by the end of the year, all of the root servers should be collecting and publishing RSSAC 02 data.

On the root server operator website, which is rootservers.org, there are pages per letter. And on each page, there is a button in the corner, as is outlined there on the screen, that will allow you to access the RSSAC 02 data. It's a YAML formatted file with all the details per day.

In addition, DNS OARC, a research group, is collecting and consolidating all of the data from all of the individual root server operators. That is not currently public. It's available to DNS OARC members. However, they're considering opening that data up. And frankly, anybody who is not an OARC member who would like to see the consolidated data, certainly just drop them a note. It would help them make the decision on whether they're going to be doing this or not.

Just one final thing, so people can see what we're talking about. That's actually what an RSSAC 02 record would look like. It's stored in per day, per metric YAML formatted files. And if you were to follow that link on the previous page, you would get a large number of files, again per metric and then per day. All the details in that YAML format.

JOE ABLEY: As I mentioned before, my name is Joe Abley. I feel like I should mention before I start talking over these small number of slides, I was certainly a member of this work party, but I was probably the person who did the least work here and deserves the least credit. So don't anybody interpret this stuff as being my effort. It's not. Dwayne and the other team members all did far more than I did.

This work party was all to do with TTLs and TTLs are a fairly obscure technical parameter in the DNS, and I appreciate that it's not as familiar to a non-technical audience as it is to people who work with the DNS every day.

So the TTL is not something that's special to the root zone. It's something that is inherent in all zones in the DNS. It's part of the base protocol. But because the root zone is a DNS zone like any other zone, all the records each have TTLs. So the question here is not anything groundbreaking protocol-wide for the DNS. It's really a question of have we chosen the correct parameters for the root zone in particular, and is there any benefit in making any changes?

So this was the scope of the work party. The current root zone TTLs, are they appropriate for today's Internet environment? There was a small change in 2014. I think [Leman] told me it was approved sometime at the end of last year to change a signature validity period, which has some impact on the choice of TTLs in the root zone. So was that change sufficient? And what impact would changing the TTL have on the DNS as a whole? That was the scope of the work party's activities.

To address those questions, Dwayne set up and divided the work into these five principle areas. So first of all, it was understanding the history of why these TTL parameter are as they are and documenting the history of it. Much of the root server system is documented only in the minds of people who were there at the time and very small amounts of it are written down. I find this worthy of – even if the

document only contained the history, I think it's worth reading. I recommend the document.

But also to try and find out whether the TTLs in the root zone, each of which correspond to a TLD, whether they match the needs of the TLD operators themselves or is there a reason there to look at changing something? And then understand how these TTLs are actually used, because without delving into the technical specifics, these TTLs are in effect instructions to resolvers that run ISPs and enterprises, campus networks, as to how those resolvers should function.

So understanding the behavior of those components in the DNS and how those might be changed in terms of how they behave by changing these TTLs, that really accommodates three and four.

Three was a survey of individual implementations of DNS resolvers and four was an actual empirical observation of what behavior do we see at the root servers to try and endorse the – so we're approaching this from two ends: what we think the resolvers do and what do we actually see them doing.

And then the last point touches on this change that was mentioned from 2014. It's a DNSSEC specific question, and it's does the fact that the root zone is now signed impose any other sorts of requirements on the TTLs or constraints or are there any changes that we should make because of DNSSEC?

So I'm certainly not going to read this through, but these are the actual values, just to make them slightly more real. The TTLs in the

protocol are expressed in seconds, but because they're large numbers of seconds, they're quite difficult for humans to read, so here they're summarized in terms of days. Next slide.

In order to answer the first question, which if you recall was are the TTLs in the root zone compatible with the expectations and the needs of TLD managers? The approach taken here was to look at the corresponding TTLs that TLD managers themselves are publishing and compare those with what's in the root zone. And we did that for the delegation sets. This is the unsigned data, the [inaudible] in the DNS. We also did it for the closest [inaudible] we could find for DNSSEC which was the DS record, the Delegation Signer record, in the root zone. And then the actual DNS key records, which is the other side of that secure delegation in the TLD zones.

I certainly won't walk you through this graph. The numbers are too small to read anyway. But the thing to notice is there are large rectangles on both sides and they're approximately the same place. What we found out was the expectations seemed to be being met. The TTLs that are being published by TLD managers correspond very closely with the TTLs that are in the root zone the majority of the time.

So this graph here, the important box again is highlighted on the right. What this really tells us is that having a TTL that's greater than one day, for the vast majority of resolvers, makes no difference whatsoever because if we give a large TTL to a resolver, most of them ignore it anyway and will cap it at a day. They won't cache records for longer than a day anyway.

So this is important because the usual tension with choosing a TTL that's appropriate is between having a small TTL, which increases traffic on your authority servers, or having a large TTL which gives you more stability in the event that you have failures.

So the question of would lowering the TTL increase the traffic on the root servers, what we discovered was largely the answer is no. We can take it down to as low as one day for all of those records and it will make no practical difference to the traffic that the root servers see.

I'll preface this with a comment. Whenever you have a technical group of people who are trying to look for problems, you desperately want to find one. We tried very hard to find some changes we could propose that would make things better and we didn't find any. It's very discouraging. It's a lot of work to do just to find out that everything is fine the way it is and nothing should be changed.

So we looked very hard in particular at DNSSEC because most of these parameters have been the way they are for 20 years or more and haven't changed at all.

We know that the DNS has changed in terms of traffic volumes and traffic patterns. Quite [inaudible] from the policy side and the economic side of it and the business side of it. So we assumed that we would find something, so we looked very carefully at DNSSEC because this is the major structural change in the root zone that's happened in the last two decades.

And we found some theoretical [inaudible] cases that in the very unlikely situation – this is being struck five times by lightning sort of territory, but we did find some areas where theoretically we could see a case in the future where we might see a problem and that encouraged us to suggest that there are some small changes that could be made to avoid that tiny possibility. I'll discuss these in more detail on the next couple of slides, but the point of point three here is that we do have a DNSSEC signed zone now, and when we do consider the changing TTLs for some reason, if we do in the future, we should certainly make sure that DNSSEC and the time is inherent in DNSSEC signatures are taken into account.

As I mentioned before, root zone TTLs appear not to matter to most clients. Again, very discouraging to people who are actually doing the work. But most clients don't really care what TTLs we publish in the root zone. We could change them to whatever we wanted and the world would stay the same.

Our principle conclusion here is that there are very few reasons we found to consider changing anything in the root zone as far as TTLs are concerned. If you read through the technical reasoning, there's a lot of quite interesting data. There's some good measurement exercises and some experiments and the data is very clearly presented in the report. It's interesting to read if you like the technology, but again, we found no smoking gun. We found no problem that we could really sink our teeth into and suggest a fix for. As it says, very few reasons. Next slide.

We had to find something, so we found something. So in the remote possibility that a root server somehow manages not to transfer a new copy of the root zone for a period of two weeks without anybody noticing, and you have a situation where a validator retrieved a particular [inaudible] response at a particular time which is within three days – there's a three-day window where if they retrieve it there, then they might have a problem, except they won't because they ignore the TTLs and [inaudible] back earlier anyway.

So this is the closest thing we could find to a problem. It's very far from an actual problem. The word problem is probably the wrong word to use, but theoretically something could happen to some validator somewhere in this situation.

Again, desperately trying to find some reason to suggest that we should make some kind of change. We identified that this minute possibility here we could even eliminate this and the document proposals. Three ways we could do it, only one of which involves changing TTLs. So this is the closest thing we found to suggesting a positive change for everything, because as it turns out, the parameters chosen in 1991 and before in fact were good parameters to choose and the root zone is fine.

I talked over that, so next slide.

The timeline here, the document is completed in draft form as being reviewed by the RSSAC caucus for I think a couple weeks now, or maybe one week. Steve's in the room somewhere. About a couple weeks now. So that process will continue to run this month.

We have some time set aside in July to change the report and change the wording, improve the clarity if we identify it could be improved in July. Then in August, it will be sent to the RSSAC exec for their consideration and hopefully publication.

UNIDENTIFIED FEMALE: Thank you, Joe. Before I turn this over to Suzanne, I just wanted to say that this particular piece of work, at the Singapore meeting we announced that we were assembling a work party to do this work. At this meeting, we have tangible findings and outcome. So this was an amazing body of work done by the caucus, so hats off to the group. Thank you. Suzanne?

SUZANNE WOOLF: Sure. The thing I like very much about having the technical reports out here and being able to report on them is that even though the technical details are obscure, we get a lot of transparency and a lot of visibility from the world into the fact that this is the work we're doing in support of users everywhere being able to use DNS and use the root server system without having to think about it at all. This is good work and it's always good to report on it.

In addition to providing technical reports on this kind of technical information, we also – another of the things that RSSAC exist to do is to engage with the larger ICANN community as an interface between the larger community and the root server operators and some of the technical considerations. So we also have been particularly recently in

connection with the IANA stewardship transition. We've done a couple of public comments on some of the planning the proposal documents and proposals and so on. We did in fact file a comment a couple of weeks ago on the CCWG request for comment on their Work Stream 1, their initial report. Next slide, please. Thanks.

It was a very brief comment, but to condense it down even further, we had some concerns that we've since discovered. We're not the only ones. We did find the proposal difficult to evaluate [inaudible], so we didn't have a lot of consensus position to provide to the CCWG Accountability group.

The principle concern that we looked at was the empowered community structures where SSAC and RSSAC in the reference model for this structure would propose to share the responsibility with some of the other SOs and ACs – all of the other SOs and ACs – for deciding whether certain extraordinary powers of the community would be triggered by particular events and so on.

We did find that some RSSAC members are uncomfortable with those mechanisms as they're currently proposed as long as RSSAC as a board-appointed committee. As an AC of ICANN, we have a great deal of autonomy over our membership, but at the end of the day, we're actually appointed by the board, so it seemed a little bit challenging for people to think about a situation where groups appointed by the board could in turn act to overturn board decisions or the membership of the board.

But more generally, and as a broader comment, we were concerned that becoming part of ICANN's decision-making processes would require changes in structure and process for us that aren't compatible with our current [inaudible] advisory committee. We've set up a structure where we think we actually function very well within our charter, which is to give advice to the community on specific topics. We're actually pretty happy with how that's working and it's not clear how we might have to change that in order to participate in these other processes and structures. So there's a certain amount of discomfort with the idea of attempting to restructure to work effectively within this proposal within the proposed structures here. When we think as an advisory body, our current structure and processes are serving us well. Next please.

The CCWG accountability chairs, they have been the busiest people on Earth this week. They [sought] meetings with most of the community groups including with us. We had a very good meeting with them yesterday. They were interested in clarification of our concerns about the implications of the proposal for us as an ICANN advisory committee.

We did have a good meeting with them yesterday and we had a couple of items to follow up. First, in the proposal, there is some very, very draft language for changing [various] of the ICANN bylaws, particularly the mission and core values section of the bylaws where the idea is to make those golden bylaws that basically fundamentally define the nature of the organization even more strongly than they do now. And they are looking for text from us as far as describing the portion of

ICANN's mission that relates to the root server system and stable and secure operation of the root server system. So we owe them some text.

And after we had described our concerns to them they have committed to making sure that those are clearly and accurately documented in their follow-up report so that all of the community understands where we're coming from with our feelings on what they're proposing.

I think that's also – even though it is very, very far from the technology concerns we have, I think that's also a successful kind of engagement for us. I guess we have a little bit of work to do in getting them some text, but it will be a good thing to have done.

UNIDENTIFIED FEMALE: Thank you, Suzanne. [Leman]?

[LEMAN]: Thank you. I'm [inaudible]. I'm going to talk a little bit about the upcoming caucus work that we have immediately in front of us, so next slide please.

We have something where we are [inaudible] taken the decision to form a work party and that's regarding naming and signing of the name of the root name servers. The root name servers have domain names themselves. These are host names, like any other host names on the network. We've identified possibilities to improve on network

traffic and security by possibly renaming them to give them other host names. These are very deep-down DNS technical details on how the names are packaged inside the DNS packet. We might be able to gain some network traffic and possibly also security, but we need to investigate that by launching a work party that gets to sit down and hash out all the details, what the effects would be.

It's certain that this would not affect the reachability of the root name servers because the root name servers are possibly the only machines on the entire Internet which you don't have to look up because the servers who actually talk to the root name servers, the resolvers, they already have the IP addresses, which is what you need to access the root name service.

So this would not affect reachability and there is actually precedence here. We have changed the names in the past. We now see that there may be cause for another change and that's what we need to investigate, whether the benefits would be bigger than the drawbacks.

It would also investigate the pros and cons with signing the root server dot-net zone. Currently, the root name servers have host names in the letter.rootservers.net, I.rootservers.net in my case.

That zone is actually not signed using DNSSEC because we've looked at this in the past and in the past we arrived at the conclusion that there's no extra security gain from signing that zone because the security is all in the content. It's all in the database records and they are already signed up in the root.

But times are changing. Things progress, so we want to take a renewed look at this to see if the environment that we're working in today as opposed to ten years ago warrant a change on the status here to actually sign this data. Next slide, please.

Another thing is a tiny little detail. The document RSSAC 002 that specifies the measurements that we do long-term to find trends that Jim Martin spoke to. When the root server operator started to implement this, they found a tiny little misspecification in a certain parameter. It turns out that there's a couple of counters that are not aligned properly, so we need to see exactly what's wrong here, exactly what do we need to measure, how we need to adjust the document so that this is actually a consistent picture, that we have consistent information in the numbers that we collect.

This work party will be charged to look at that and produce a revised version of RSSAC 002. So this is a very small thing, actually. Next slide, please.

We've also identified some potential future work, and this is work where we see that maybe RSSAC isn't the right place to conduct this work. So we're discussing and looking at interacting with other bodies in the DNS arena to help us or give input on this.

The first one is very much within our remit. It's improving information about the root server system and making that information more accessible because we see when we talk with people and when we receive questions that there's lots of misconceptions out there about what the root server system is and how it works and how it interacts

with other bodies. So we need obviously to provide more information and try to push this out in order to give people a correct picture of what it is and how it works.

We also talked about creating a test bed to validate root server conformance, the RSSAC 001 document. Service expectations has a sibling document which is about to be published by the IETF (the Internet Engineering Task Force) [inaudible] body.

It turns out when you start to look at the requirements on the root name server that you can divide it into different parts. One is the protocol requirements. Which parts of the DNS protocol should a root name server adhere to and respond to? The other one is what are the service levels that you should expect, the qualities – or rather security implementation, response times, things that you can measure in that respect.

It would be good to have some central way or some unified way to test the root server system and the various components of it and make sure that everything actually follows these expectations and requirements. But we see that RSSAC is probably not the right place to do that because we are comprised mostly by the root [server] operators and it would be better if someone else actually checks that we follow these expectations and requirements.

Another thing we identified is to look at whether to expand the so-called DITL measures. DITL is short for Day In The Life of the Internet. Every year the root server operators participate in a large effort where all the incoming DNS queries to all root servers all over the world are

collected – at least that's the theory – for two days. So we have a sliding 24-hour window that we can look at.

They are then stored in a gigantic database and can be used for research efforts. People who want to look at trends and qualities or queries. It has actually been put to very good use already. It helped out figuring out, for instance, which new gTLDs were unsuitable to delegate because there are so many queries for them already that there's obviously a larger user group out there who would suffer if you had real domain names delegated with those labels.

We are thinking around the problem of [inaudible]. The root servers are hit by various types of attacks, like any other server on the Internet. And we sometimes see rather sudden traffic spikes that we would like to be able to mitigate the effects of the spikes, but in order to understand what's going on, we need to analyze the traffic. But these spikes happen at irregular intervals, so we don't know really when to collect the information. So we might have to expand the DITL measurements to longer periods of time and collect data over longer period of time to catch the traffic when it happens.

But this is just a very basic idea. This is a seed of an idea. We don't have the resources to actually do the database crunching and the storing of all this, so we need to reach out to other bodies to help us do that. So that's what we've been looking at. Next slide, please. Yes, you can go for the next slide again.

We have too little interaction with other bodies. We try to reach out, but there seems to be a lack of knowledge on both parts – on our part

on how to reach out to the various constituencies of ICANN and other groups, and also we see when we talk to other groups that the understanding of what the root server system is and what RSSAC works and how RSSAC works inside ICANN is somewhat limited.

For those of you who are here and who are typically not RSSAC members already, we are seeking your help. You're obviously somewhat interested in RSSAC because you are here. Do you feel like you can find the information that you would like to have, and if you can't, how can we improve on it? Are you aware of the various ways you can interact with RSSAC? I would encourage participation here. Please give us your feedback.

I guess we can extend this also to questions regarding the presentations that you've seen here so far. Please, are there any comments, any questions? Complete silence.

UNIDENTIFIED MALE: Hi, my name is [inaudible]. I am a fellow and I'm from Dominican Republic and I am from ISOC. First of all, I would like to – sorry about [inaudible] talking a little bit. First of all, I would like to thank to Mr. [inaudible] who was with us today this morning. He was explaining many things.

For my part, I would like just to say from our community, we just – maybe there are, let's say, basic concepts about this and how we can benefit from the knowledge and from knowing more about the project and all those things. How can we get more involved?

For example, today we have another topic [inaudible] being deployed and also we are talking a lot about it. It is about ESP. For example, I would like to know if there is some relation with it, because we are watching that. When we have, for example, ESP in the region, it is good for community and its cost for Internet broadband. But besides that, it could be something like making regions a part and things like that. I will find out, for example, if there is no [inaudible] with root servers in that condition.

UNIDENTIFIED MALE: I guess I was going to ask a clarifying question. Are you asking about root server deployment; in particular, IXPs?

UNIDENTIFIED MALE: I mean, are IXPs being deployed around the world? If there were IXP around the world impact the root server?

UNIDENTIFIED MALE: Excuse me, are you saying ESP or ISP?

UNIDENTIFIED MALE: IXP. It is Interchange Point.

UNIDENTIFIED MALE: Oh, IXP, okay. I am not a root server operator, so let's make that clear. I used to be one. I no longer play one on television. There was a lot of interaction between people who were promoting connectivity,

especially in developing parts of the world where you have two things to fix. One of them is internal connectivity between different networks in a particular region. That is the part that the IXP addresses directly. How do we exchange packets directly? How do we grow our traffic between us and improve the number of services available on the Internet? IXPs address that very directly.

But what also tends to follow with some of these projects – and I'm sure some of my colleagues here who do run root servers can comment more – is bringing resources and services to the IXP provides more reasons for people to connect and provide local content so you don't have to go to other countries or external providers in order to reach it.

So I think they go together very well, and I think root servers are a very good example of infrastructure that can be brought to an IXP and strengthen the arguments for individual providers to connect and exchange traffic.

UNIDENTIFIED MALE: I will pick up on that and speak for NetNode who operates one of them. I know that several other operators – I [won't] say all, but several others are in a similar position. We try to follow the development on the exchange point side because it's a very important channel for us to reach out to local communities. But there are a couple infrastructure things that are needed in order to – for a root name server to function well at an exchange point.

Two of the more important ones is that we need to be able to reach the server from the control center. In our case, that's [inaudible] Sweden. But over the Internet, we must be able to load it with new data. As the zone file are updated, we need to transfer data to the machines. That's one problem.

The other one is that we need to find a stable environment. And I'm not talking about the physical environment with power and cooling, but a stable social environment where we can connect with the local engineers on the side so we can have someone to talk to if we have problems, if we need to replace a hard drive or ship new parts or just press the reset button when something goes wrong – to have that interaction, to find the persons to interact with is often a problem.

And, of course, as always, to find a financial solution so that we can pay for the hardware to ship it there. That's often a rather small problem. There are several solutions for that. The various root server operators have different models for how to solve that. That's not a big issue.

But getting the transit traffic and getting the social environment sometimes poses challenges because some of these exchange points, good ideas as they may be. They may not survive in the long run. After a year or two, the person who had the energy and started everything leaves to do something else and it kind of decays. Then the root server is no longer very effective there. At worst, it may not even help the local community, but pose a problem instead.

There are some problems involved, but it's also a very important channel for the root server operators to reach out and provide good connectivity in the local environment.

Any questions?

UNIDENTIFIED MALE: I just want to extend what [Leman] had said. My group operates F-root, and we have over 50 nodes out in [inaudible] in six of the seven continents at this point. What we find is that it dramatically reduces latency for the end users. It also reduces costs for the ISPs, because they don't have to have this traffic go out over the transit links.

So while [Lemon's] points are all entirely accurate, as a general rule, putting root nodes into IXs is almost always a win. If you're aware of any IXs that don't have root servers in them and you have any influence over, please send them to us.

UNIDENTIFIED FEMALE: Just to add to that, there is a website root-servers.org, which has a great deal of information about the various root server operators, which exchange points they operate in. If you want to partner with them, there's information how to get in touch with people and how to start the process [Leman] described. There's actually a fair amount of good information there. Talk to people. Get them to work together.

UNIDENTIFIED MALE: Any more questions or comments? Because if not, then I would like to say thank you all for coming here. Do come and talk to us in the hallways. We are very happy to talk to you. We would like to exchange information. We are happy to tell you what we do, who we are, how we interact, how it all works. Please come and talk to us. We'll really appreciate that. Thank you very much.

[END OF TRANSCRIPTION]