Julie Hedlund: Hello. Good morning. This is a double session. We are beginning this morning with the internationalizing WHOIS preliminary approaches for discussion with Jeremy Hitchcock, co-chair of the Internationalized Registration Data Working Group. Welcome, everyone. A little bit of logistics. We'll go ahead and start with the presentation, and then we'll allow some time after that for questions before we go to the next presentation and what we'll do is ask if there are questions in the room here. And we'll ask that you use your microphones when you respond, as well as please introduce yourself before you speak.

> Please speak slowly and clearly, because this session is being livetranscribed. So it's very important for the scribes to be able to hear you. So we'll ask for questions here in the audience. And we also have a live chat room, and we will ask for questions in the chat room as well. Thank you, everyone, and welcome to Edmon Chung, the co-chair also of the Internationalized Registration Data Working Group. And at this point, I would like to turn over the session to Jeremy Hitchcock. Welcome, Jeremy.

Jeremy Hitchcock: Good morning. I'm Jeremy Hitchcock, Co-Chairing the IRD working group, along with Edmon. Kind of as a general bit of background, this was a working group that was spun off to look at what type of impacts, requirements for WHOIS we might have in a -- in internationalized context. One more? So the group started in December of 2009, and since then, we've been doing biweekly meetings to scope out the impacts of internationalized content in WHOIS. We've been doing them in a couple of different time zones to ensure that we have participation from different geographic regions.

> So we're going to give you an update on where our current conversations have led us and some possible approaches that the working group is going to explore over the next few months, and the hope of producing a report by November of this year. The general scope for WHOIS -- and WHOIS has had many studies, and many times has been looked at. The particular thing that we're looking at is, under an internationalized context, especially in the advent of many new gTLDs, the question is, what sort of internationalized content in WHOIS should be displayed and how it should be displayed.

Right now in the RAA, it's simply listed as name and date, and there's really no specification that's given for internationalized content. And the

port-43 WHOIS service that's discussed in the RAA refers to a number of RFCs that don't really discuss what the requirements are or what sort of encoding for any sort of internationalized content. So this working group is looking at what the -- what impacts or what sort of things we might require from WHOIS to display internationalized characters.

So this is a Joint Working Group with SSAC and the GNSO. We have representation from a number of ccTLDs and different geographies, different implied language scripts. And the influence from the ccTLDs in the sense of what they have done in the past for looking at internationalized content has also been helpful, because, obviously, they've had to deal with the issue of presenting local script in the WHOIS protocol, which doesn't really extend nicely to internationalized characters. The three main approaches that we're looking at is the general service requirement.

So what should actually go in and come out of WHOIS in the sense of what is the WHOIS service intended for in context with internationalized content, looking at the localized user experience. In this sense, it is how does a consumer of the WHOIS service with internationalized characters, especially with cross-language scripts, experience WHOIS. And we're also looking at what are the input side and the display side of WHOIS, and looking at potential alternatives that might be useful. So the basic premise of the general service requirement is what capabilities are needed for the WHOIS service in the IDN environment.

For domains, we have Punycode, which is a good representation for internationalized characters in domain names. But that really doesn't cover any sort of content information. The screen shot -- can't exactly see, but if you download the slides, you'll be able to see the exact specifics. Port-43 clients as a potential service requirement must be able to display the U-label or the A-label. And so when we're discussing U-labels, we're talking about the Unicode format of the IDN label. And then the A-label is the ASCII form of the IDN.

And so in the first box, you see a – the actual Unicode characters, a bunch of Chinese characters, and then the second one is the A-label representation, which starts with the xn--. So it's an ASCII representation that ultimately can get encoded. But the client is required to do any sort of representation to show that to the user in kind of their native script. WHOIS clients must be able to display the results of queries for domain names, and the other alternative is the bundled representations of a single A or U-label query should be returned. So those are the three potential options that we're exploring. A couple other things that have come up is that the WHOIS protocol has no mechanism for indicating a preferred character set either to use for the input or the display.

Some potential things that we've been thinking about is looking at was actually given in the query string from WHOIS because WHOIS is basically a single-key lookup. So maybe something is returned differently if it's provided and the query string is an A-label or a U-label. And that might dictate some approaches that would have to be explored with IETF as a standard track for the RFCs that relate to WHOIS. Next slide.

The localized user experience -- so we're looking at what are the requirements to accommodate users who have to submit registration data in familiar character scripts. If you are a local user in China looking up a domain in Chinese that is one particular flavor that can -- has been explored a bit more. But if you are a Cyrillic character set as your language local script, how do you look at something that is actually Chinese? And that's something that we're looking at a little bit in terms of kind of an education that's somewhat interesting that hasn't been explored completely yet.

So the specific data for WHOIS that's discussed in the RAA includes domain names and name servers. And those two have been somewhat explored through the Punycode. But there's other contact information that's provided. And that includes the sponsor and registrar, both the name and contact information for that registrar. There's telephone and fax information. Their e-mail addresses, dates, registration statuses, and then the entity names, the registrant, the admin contact, technical contact, and the corresponding postal addresses. And so those are the places of contact information that we're looking at in the localization.

So for domain names, the WHOIS service, even though we have Punycode that provides an A-label, one possible approach might be to suggest including the Unicode characters in the WHOIS results, so maybe we provide back both variants. For the name servers, it's possible that some will publish their name servers in IDN. And the question there is, is this the same as far as A-label and U-label representation, potentially providing both as it's available. For the sponsoring registrar, this has been a topic that has led us to a couple of different options. The sponsoring registrar is usually a piece of information that is interesting to consumers of the WHOIS protocol in the sense of law enforcement or business uses or intellectual property. There's a vast list of reasons why somebody would want to get a hold of the registrar. The working group has been discussing whether or not to have this information always available in ASCII. And then to the extent, consistent with the registration -- registrar accreditation process, also make it available in the local script.

For telephone and fax numbers, the ITU E. 123 standard has a nice notation for phone numbers and fax numbers. And I think the working group is fairly happy with using that as the representation for telephone and fax data. Dates is a piece of internationalized content that can be represented in many different ways. And I think that the working group felt that the EPP representation for dates, which is denoted as year, dash, month, dash, day, I believe it's a Z -- or, no, a T. I can't remember which character which separates the time stamp. And the time zones are also covered in that.

And then that's the -- the approach that we have taken for expiration dates, creation dates, and et cetera. The registration status is something we have a few different options that we haven't completely discussed yet, so the registration status includes client hold, client delete, prohibited, update prohibited. I mean, this is how somebody can determine whether a domain is locked. So the options include leaving it as ASCII-7, always publish the actual EPP status code, identify in a more easily representation for the character set in which the other contact or registrar information is given, or publish easily understood representation in mandatory and local character sets.

And, of course, any of these options can be kind of mixed in terms of coming up with a single way that the working group goes forward. And then the three approaches for displaying entity names, which is registrant name, administrative contacts, technical contact, and the corresponding postal addresses. So approach one is registrants must be in a present language. The current standard is U.S. ASCII. So registrars are responsible for providing the ASCII version of the data, whether or not there's internationalized data that's associated with it.

Approach two, registrants submit in local scripts and registrars provide a point of contact. And so the -- this is similar to approach one with the requirement that the registrants provide contact data in their local script, so it's useful for their local community. But registrars have to be a point of contact to deal with translation issues in terms of getting a hold of the different contacts that go with a domain name. And then approach 3 is registrants submit in local scripts and registrars provide some sort of transliteration.

Registrants may not be necessarily savvy with how to transliterate to a U.S. ASCII or into an A-label. And the working group has been looking at libraries sort of do this in an automated fashion. We're certainly sensitive to the lack of accuracy in the sense of always being able to translate different character sets into ASCII. Some just general discussion, and this is where we would appreciate some feedback and commentary, because we're looking at, basically, functionality that's not included in the current port-43 WHOIS. Some backward compatibility issues are certainly coming up.

And as a little bit of background, different ccTLDs have approached this problem in different ways. One of the ways is to look at what the query string is. Another way is to allow for somebody querying WHOIS to provide a character set they're looking for. So they might request ASCII or they might request UTF-8. So some issues that we would like to discuss is, you know, complete compatibility with existing port-43 requests and response in ASCII only; some sort of enhanced port-43 request allowing domain names in U-label and A-label format; enhanced port-43 requested script code; enhanced port-43 response, allowing ASCII and UTF-8, shifting to another port, replacing port-43, or, as the RAA requires, there's also a Web-based name query tool.

And in the HTP specifications, there's little bit more framework for languages and accepted languages. The HTP headers include an accept language set which says what's -- it's intended for which language scripts the browser is able to support and what the user is able to understand in terms of characters. And that might be a potential way to move forward. So where we are is, we've basically kind of set the groundwork for what exists currently. We've looked at what sort of impacts those might have and we're starting to discuss those and we're considering what impacts and benefits to users and stakeholders have to those different approaches, looking to recommend a way forward and produce a report, again, by November of this year. So those are the, that's an update of where we are and at this point, happy to take any questions and look for some discussion on this.

Julie Hedlund: No, it's not working. Try a, try a different --

Jeremy Hitchcock: There you go.

- Alex Mayover: Alex Mayover from Nikto AD. I have to admit I am completely new to the topic. The IETF has produced, actually, a -- sort of a replacement protocol for WHOIS that's called IRIS that's not on the list. Could you elaborate a little bit why this is not on that list of options?
- Jeremy Hitchcock: The two things that I would point to in terms of IRIS, it certainly has been part of our discussion. One has been the workload and probably polarity of WHOIS, and, you know, that's certainly one potential way that the working group may say is that replace WHOIS, get rid of it, and put something else in. I think a drop in replacement for WHOIS using something as IRIS that certainly has more policy implications.

And so whether or not people will look at the work that we're doing and say that that constitutes a reason to replace WHOIS is certainly one way. Kind of a more generalized technical reason why not to is that right now, there's only one implementation for IRIS, and so in kind of the -- the way in which a protocol or a -- or an implementation of a protocol can be used, there probably should be more than just one implementation that exists.

And so if that's something that becomes a serious option or a serious consideration, that might be something that would have a stronger implication in the sense of wanting to see more versions of IRIS or crisp that exist out in the wild.

- Edmon Chung: Just to add to that, I guess, there is quite a number of works done on WHOIS right now. And part of the thinking is not to preempt those discussions. And I think that should be an important consideration as well. And a lot of the work done here, regardless of IRIS, needs to be addressed, like -- 'cause even with IRIS, you still have the issue of do you want a must-have language or -- you know, those kinds of things. So -
- Julie Hedlund: That was Edmon Chung who was just speaking.

Edmon Chung: Sorry.

Jim Galvin: And just speaking as a member of the working group, my view about IRIS is that as compared to what this group is doing, this is really more about the issues and identifying them and specifying the requirements for the solution. I mean, IRIS may very well be one of the ultimate things which meets what it is this group is looking at in defining the -- what it means to have internationalized registration data in the WHOIS service requirements. So it's not that it's not part of the discussions. It's just not time to -- to make that choice.

Greg Aaron: Hi, Jeremy, Greg Aaron from Affilias. When looking at some of the potential technical solutions, one of the things that I hope will come up are some of the complexities and kind of not worst-case scenarios, but most complex scenarios that might come up. I was looking at all the variants that are produced when one implements, for instance, Indic languages. India is a place that has, I forget, either 15 or 22 official languages.

It produces a lot of variants. And it might not be possible to produce variants in WHOIS because there could be hundreds or thousands for one domain name. And then we get into an issue of whether that blows SLAs for WHOIS provision by registry. So I was wondering if the group had been thinking about all of the complexities around that.

Jeremy Hitchcock: Yes. Generally, yes, we've been looking at especially the complexities dealing with transliteration and providing some sort of translation process. Google translate gets you some of the way. But it's not something that necessarily would want to rely on completely. And it would probably be a pretty expensive burden for everything to be translated into a multitude of scripts. So that's something that we're -- that we're looking at.

I think one of the pieces of conversation that we've had is that what's the ability to drop a piece of mail in a local Postal Service and what's the extent possible to have that piece of mail delivered to another country, even though it's a different language script. And we've been using a similar metaphor, I think, and applying it in different ways as far as trying to think about cross-script compatibility or cross-scriptic access. But there are a lot of variants out there. And trying to provide interoperability between them all is not an easy thing.

Julie Hedlund:	Are there more questions in the room? Do we have questions in the chat room? There are no questions in the chat room. Anything anybody wants to add before we move to the next presentation?
Jeremy Hitchcock:	Could I make a general request for anyone who is interested or has interesting cases that WHOIS in an internationalized world doesn't support, send them to the working group so we be sure we consider all the options and alternatives and try to think about all those interesting edge cases to support the WHOIS interoperability as best we can.
Julie Hedlund:	Thank you very much, Jeremy. Thank you, Edmon. And I think we'd like to go ahead and move to our next presentation. Our present I'm sorry, our presenter will be Dave Piscitello. And it will take a moment to get the next presentation up online. Okay. We'll go ahead and get started. We have the slide deck ready. And I will turn this over to Dave Piscitello, senior security technologist, ICANN.
Dave Piscitello:	Thank you, Julie. And good morning or good whatever it is wherever you are. I'm going to to bring you up to speed with some work that primarily Steve Sheng at ICANN has been doing in terms of investigating requirements for future WHOIS services. To provide you with a little bit of background, the GNSO Council asked that policy staff compile a comprehensive set of requirements for WHOIS service policy tools.
	And the goal was to get a very, very comprehensive list of features and service elements that would not only reflect features that could fill in needed gaps from known deficiencies, but also to include any possible requirements that had yet been overlooked or that might have been studied in the past and might now be relevant again. The second part of this study, which is still a work in progress, is to develop a straw man proposal based on the requirements and based on some feedback from the community.
	Before I start, I want to make certain that we put the context of the word "requirements" in the right context. We are when I say "we," I mean support staff are looking at current features that are identified as needing improvement, features that support various past policy proposals, and also features that have been recommended by various ICANN service organizations, advisory committees, and the community at large. We're not gathering policy requirements. This is technical. This is looking at the protocol itself, what you know, the services that are built upon the

protocol and what one might need to do to enhance the protocol or, in the future, perhaps, supplant the protocol with something that would allow us to have a richer user experience for WHOIS.

The status of the report is that a draft report was released in March of 2010 and sent to the SOs and ACs for their input. During the months of April and May, Steve Sheng conducted webinars with reasonably good attendance in both to share the original draft report. We then received some input from the Registry Stakeholders Group, the GNSO, the ALAC, and several technical experts. And Steve incorporated those comments and released a draft final report at the end of May.

I'd like to go briefly through some of the comments highlighting some of what I believe is the substantially positive feedback that we received on the set of requirements that we identified. The ALAC did support all the requirements. They believe that there was a consensus in the community on these. That's their expression. The registry -- RySG was happy that we had gone through the effort of collecting the requirements and believes that it's an excellent addition -- excellent basis for additional definition WHOIS service requirements.

Before we proceed and look at some of the details of what's in the report, I'd like to make certain that we're on the same page when we talk about WHOIS service. The WHOIS service is actually comprised of several components. Obviously, there's the user. But beyond the user, there is a client that he or she uses to access registration data using the WHOIS protocol. There are many forms of such clients. One is a text client that actually connects using TCP to port-43 at a server, and then places text queries and receives text responses.

There are various forms of Web-based clients that have form submission pages and display mechanisms to provide not only the registration data that are collected by registrars, but also some additional data that registrars may wish to offer, and registries as well. And then there are also clients that are automation. By automation, I mean applications that actually go out and grab WHOIS records and process them in some way, either parsing them to get specific information out of the record or capturing large numbers of WHOIS records and trying to examine or determine some patterns in the WHOIS records that they have captured. And often those clients are used to do some form of monitoring or analysis. They are used quite often by people who are investigating phishing or spam or some other malicious activity. On the opposite end of the WHOIS connection from the client is the server, and the server is actually processing the queries and examining or collecting the data that the query requests and returning it in the form of a response. And then last but perhaps most importantly is the registration data themselves, and the data are collected at registration time and they are also updated and maintained as the registrant and registrar need to do so.

As in many cases, terminology has meanings among different groups and constituencies. In the particular case of the terminology that we identified at the outset of the report, the ALAC had some disagreement as to whether or not a Web-based client should be considered a client at all. Their argument is that a Web-based client doesn't suffer the same limitation as text-based client and can handle some of the service requirements that we mention later on in the report, such as authentication, internationalization, and anti-abuse features.

So now our initial compilation, we came up with this laundry list of services and features that would be of some benefit and would complement the existing set of services or improve them. Since I am going to go through each of these one by one, I am not going to read them off this page. One of the things that is very challenging for anyone who uses WHOIS service on a regular basis is actually finding the authoritative WHOIS servers for specific registries or specific registrations.

So one value add that the community in particular has indicated would be particularly helpful is to allow the community to migrate away from the way that they currently grab and identify WHOIS servers. So a list of domain names and IP addresses of all the authoritative WHOIS servers made available publicly would probably server users better than the way they do so today which is to use clients heuristics or preconfigured tables, or simply the fact that you know if you need to get an IP or a WHOIS record from a specific registrar, that you go to WHOIS.registrar.com.

So, one of the first technical requirements in the report is to provide a publicly accessible and machine parsaible list of domain names or IP allocations of WHOIS servers. And this could be something ICANN could maintain in the same way it maintenance the list of the accredited registrars or any other list at its ICANN site. It would also be valuable if

such a list would be made available in a machine format so that someone who is actually building automation could grab that list and incorporate it into an application. Another area that has evolved over time to become important is the ability to structure queries in such a manner that you can actually appreciate the format of the query data.

So, for example, if you go to the Internet registry world to query an Autonomous System Number, you have specific elements in a WHOIS command line that you must add, flag according to the Unix jargon, to actually get an IPv6 address. As an example, the ARIN WHOIS interface, you would say WHOIS minus H, whois.arin.net, lower case a, and then number 6. We see the same sort of patterns emerging and the same deviations from any standard or convention in the early deployment of WHOIS implementations that support IDN responses.

So for example if you go to dot DK to query and get back an IDN domain name, you would have to add hyphen hyphen charset equals Latin minus 1. If you go to the dot JP WHOIS, you would add a /e. If you go to dot DE, you would add minus C UTF-8. So you can see this is not only problematic for automation but very problematic for users. You have to understand what it is exactly you need to submit. If it's three forms on conventions, that's not so bad but it could grow to hundreds if everyone chooses to do so differently.

And that makes developing WHOIS clients very clumsy, understanding how to incorporate all of these in every one of them very clumsy. So the second recommended requirement is that we define a standard query structure that allows all clients and all gTLD registries and registrars to essentially use the same convention when they ask for specific WHOIS data. Talking a little bit more about a standard set of query capabilities, the past GNSO and SSAC reports have called for not only you a baseline of query capabilities but the ability to expand query capacities beyond simply asking for a domain name.

Some registries actually offer certain expanded search capabilities today. You can query for the sponsoring registrar from certain registries, you can query for specific other kinds of data. One of the additional requirements we have talked about then is the ability to permit users to submit any registration data element – by "element," I mean something like a contact name or a contact address or a telephone number or a domain or name server as the query argument. So as opposed to only being able to say WHOIS example.com, you could not only ask for the domain but you could ask for the sponsoring registrar of the domain or the registrant of the domain.

This is not a straightforward implementation, and the RySG did point out that this does pose not only significant technical issues but there may also be some issues relate to go Service Level Agreements that registries are obliged to meet with their contracts with ICANN. It also could be a way to farm information that could be used in a malicious manner. In particular, it could be used to farm personal information. So ALAC and technical experts are concerned about the privacy implications of such a feature.

So while we are recommending these things as something that might be valuable from a technical perspective to be able to do, it's very clear that having the technical platform to do so then does open a policy discussion. Moving on to responses, one of the -- another of the problems that many WHOIS users and, in particular, people who develop applications or automation experience is that depending on which WHOIS server you query, you get back a different formulation or composition of registration data and labels. And this makes it extremely challenging to try to examine and correlate information across multiple registration records.

So one of the value adds that might be particularly important to consider is having a standardized format for responses in WHOIS queries. So we have actually two recommendations in the report. The first is to define a standard data structure for WHOIS responses. And then the second tries to impose a little bit more formality on the data that are returned. And so the data structure should provide for the correct identification, syntax and semantics of each data element. So if, for example, the data element is a phone number, then -- as we saw in Jeremy's presentation, it might be useful to choose an internationally recognized format.

If the particular query argument is for a status code from EPP, then it might be useful to have a single convention for how we represent that to users. We did get some positive feedback on this comment from ALAC who said that this would allow for easier localization of client software. Lastly, among the sort of interaction between client and server, if there is a problem in processing a query today, the error message that you receive back from a WHOIS server is going to be very server dependent. They handle errors differently. Some have a longer taxonomy of errors than others. And so the lack of a standard error message actually not only creates ambiguity and confusion for users, but it also makes developing applications somewhat challenging as well.

So another recommendation in the report is that we define a standard set of error messages and a standard handling of error conditions by servers. Some of the errors that weeds consider appropriate for standardization would include some sort of standard signal that the number of queries the client has issued exceeds the limit that the server has imposed. Many registrars and registries will restrict the number of queries that you can place from any given IP address or from any given client to prevent the automated harvesting of WHOIS information.

Another is to have a standard error message for "no record found." And another would be to basically say I can't process the query you submitted. We are going to move away from the client/server interaction and talk more specifically now about the quality of domain registration data. And by quality, when we are talking about databases, we often have three characteristics that we try to examine when we look at data. Are the data accurate? Are the data useful or relevant? And are the data -- or the collected data current?

So we're looking at integrity, currency, and utility. As soon as we get into the issue of looking at data quality, we've run head-on into the issue of trying to maintain or improve WHOIS accuracy. And WHOIS accuracy is actually affected by a number of different variables. The report goes into some details explaining what these are, but at a high level, the general issues of trying to maintain WHOIS accuracy have to do with privacy considerations. And in privacy considerations, there are people who don't want to disclose their personal information, even though they are obliged to have contact information as part of a registration process and so they will actually populate the WHOIS record or their registration record with false information.

Obviously, someone who is going to take a domain and use it for some malicious activity is very, very interested in stealth, and so that party will engage in some intentional deception. They will falsify the record because they want to lead law enforcement or some other responder down the wrong trail. We also have very little in the way of corroboration of submitted data performed by most registrars. There isn't any confirmation mechanism, there aren't very many registrars who actually will go out and verify that the e-mail address, for example, that is submitted is a working e-mail address and the like. And so there's a lot of room for improvement in those areas, but there's a lot of complexity in making those improvements as well.

And then of course there's also user error. And there are some people to are first-time registrants and they actually make a mistake when they enter their registration data and they don't go back and change it. While we were looking at registration data, one of the things that we did was sit back and sort of study what we have collected over the past 20 years to try to understand how relevant and useful the data that we collect are today. There are fields that many people do not populate any longer because they don't use the technology.

One of the most obvious of those is facsimile. A very, very large proportion of registration records don't have a facsimile entry. So one question would be whether we need to deprecate that field and perhaps consider some alternative field that is more popularly used today. What people use today in place of fax, especially for a short message, is a text or SMS. One could even envision that there might be some people who would like to be tweeted. So what our conclusion was in this area was the data model that we're considering or we would consider for future WHOIS service should be extensible and should be modifiable.

We should be able to add or delete elements as they become relevant or irrelevant. So the two recommendations that come from this particular part of the study are that we would like to see the community adopt a structured data model for WHOIS that is extensible and changeable. And one of the fields that we actually thought was particularly important, given how important registration data are in many, many communities is to add a time stamp that would show when the field was last verified or updated.

So this would be very similar to any normal database activity that illustrates when transactions were performed or executed or when data were touched. We did receive some comments from the ALAC in this area. And one of the comments that was particularly useful was that the ALAC pointed out that structured data would allow us to use internationally agreed standards for display of such fields as postal addresses and phone numbers. They also acknowledge that the machineparsable output from WHOIS servers would be useful or beneficial for legitimate uses of WHOIS information. And they did acknowledge that when we make things easy for legitimate users, it's always hard not to make them easy for people with malicious intent. The caution that ALAC expressed is that it would be particularly important in the policy deliberation aspect of whatever future work we would do in WHOIS should consider mechanisms and policies to prevent large scale harvesting of data for malicious use.

We had considered the internationalization data issues briefly while we were doing our initial analysis and then actually realized that it would be very, very useful to simply defer to the Internationalized Registration Data Working Group and work with them and incorporate their recommendations and requirements into the general set of service requirements when that working group actually concluded them. One of the things that the current WHOIS service clearly lacks, being such a very, very simple protocol, and especially in the text incarnation of a client, is the ability to do identity assertion, credentialing or some sort of data authentication.

All these are security features that would be particularly beneficial and could, in fact, be some of the mechanisms we could use to prevent malicious use or abuse of the data or harvesting and mining. So it would be really nice if what we could do is provide mechanisms to distinguish natural persons from artificial persons. It would be nice if we could protect or discourage harvesting and mining of personal identifying information. Once we express those desires, the difficult part is to try to do so with the protocol that we currently have.

So obviously it would be important for us to start to consider something on the order of the kind of security framework that you see in enterprises or other organizations who have to routinely deal with data that may include personally identifying information as well as other data that is potentially sensitive or should not be tampered with. So three areas that we recommend and discuss in the report are the need to provide a technical means of performing authentication, technical means of providing access control, and the ability to audit.

When I talk about WHOIS security frameworks in those three areas, it's probably useful to sort of dig down a little deeper and explain what we mean. An authentication framework is not simply being able to identify the user and challenge him with a password. It actually allows you to build a framework that accommodates both an anonymous access as well

as verification of identities and also allows for a choice of authentication methods and the kinds of credentials that you would request or require of a user so that you could have that user demonstrate that he is who he claims to be.

The authorization framework is really what most people call access controls. And access controls can be very coarse, meaning once you are authenticated, you can see anything in a database or all the registration data, or they could be what we call granular, meaning that there are individual data objects, such as a contact name or a contact address, and each of those data objects might have a unique permission. So it is possible and it is done in many other databases to grant the permissions for one party or one group to see certain information and to restrict the viewing of that information for someone who chooses to access the data anonymously.

So having a strong authentication framework is actually very valuable for trying to achieve some of the very, very challenging and tricky policy considerations that we have to try to meet in this community. And the last framework is an auditing framework. One of the things that is very important to understand is how people are using WHOIS and who is using it, and understand how we can trend and see what information is most often accessed, who has accessed it. And so having a set of metrics that we could use to measure WHOIS access would be fairly important if we want to actually manage the service more effectively than we do today.

Some of the comments we received here and I think one of the most important ones from ALAC indicates that the authentication framework is a fundamental prerequisite to allow for the protection of individuals. I think that pretty much says it all. Another requirement that actually came from the Security and Stability Advisory Committee is that registrars and registries should provide and publish abuse point-of-contact information as an element of a domain registration record. So this is an example of information that we would like to see added to a WHOIS record, in addition to the sponsoring registrar.

The RySG had some comments and some additional suggested requirements. They thought that it would be very valuable to consider means of ensuring consistency of data between registries and registrars. They acknowledge that accommodating the privacy services is important, but they would like to make certain that it does so in a manner that

	effectively provides access to information for those who need it. And they want to make certain that the requirements that are proposed are done in such a fashion that they mitigate impacts to SLAs or to the underlying protocols that the registries and registrars use to do provisioning; in particular, the Extensible Provisioning Protocol, or EPP.
	The RySG actually had a suggestion for what our next steps would be, which is helpful. They wanted to have staff go back, take a look at the requirements, and then try to understand which of the proposed requirements would require or involve the examination of or perhaps modification of Internet standards. And so we will do this, and we will actually try to make certain that if there is something that we're proposing that would require a change to a standard that we take measures to actually initiate those changes.
	The ALAC and other experts actually wanted to make certain that the community gets an ample opportunity to discuss the services and to have a role in the decision of which of these requirements would move forward. ALAC actually is, I think, relatively enthusiastic about the work, and would like to see a road map and a timeline with milestones that would describe the implementation plan for the requirements that are chosen by the community.
	And ALAC and some of the experts would actually like to see some assurances that there's backwards compatibility and a transition plan once we have decided what we want to implement, we have the policies in place, and we begin to make the implementation change. That concludes the slide presentation. I am happy to entertain any questions that people here might have or that people in the chat might have. Thank you very much.
Julie Hedlund:	Thank you, Dave. Do we have any questions in the room?
Jason Poulos:	Hi, Jason Poulos here. With regards to providing abusive point of contact, are you talking about abuse by the registrant or abuse of the domain by someone else and you're notifying the registrant of the abuse?
Dave Piscitello:	What we mean is the contact information that would get you to a party in the registrar who has the has the knowledge and the authority to take action on an abuse complaint.

Julie Hedlund:	Additional questions? Are there any questions in the chat room? No questions in the chat room. I see we have one more question.
Unknown Speaker:	No more questions.
Julie Hedlund:	I see we have one more question.
Jason Poulos:	Yes, I have a couple more questions, Jason Poulos again. Has there been any need to see historical data. You were talking about time stamps, R6BS, time stamp of last update for a field. Is there a requirement to see previous data for that domain name?
Dave Piscitello:	I'm I'm having a senior moment. I know that at one point, we had discussed a what is characterized as a WHOWAS service. And I I thought it was still in the document. And I honestly now cannot remember whether it made it into the final draft. It was considered I personally consider and this is not a staff position, but or a community position, but I personally consider historical data extremely relevant.
	And it's very evident that there are other people who actually consider it relevant and valuable enough to charge subscription service access to such data. So there are companies like Domain Tools who provide historical WHOIS. You have to have gold or a platinum membership or whatever to get access to that. So -
Rafik Dammak:	Rafik Dammak. So you mean access to historical data?
Dave Piscitello:	Historical in other words, the chronology of registrations relating related to a specific domain.
Rafik Dammak:	But you know that such historical data maybe have a problem of privacy?
Dave Piscitello:	Absolutely. I mean, there are issues that you have to consider when you are going to share that kind of data. But I don't think that they're any different than the issues you consider, you know, when any of the contact information is personally identifying.
Rafik Dammak:	Rafik again, but how to say I can understand that to access to some data, current data. But to access to previous data, especially I'm not sure if domain names changed from an owner to another owner, so what is the relevance to access to such data, and maybe defer to some context?

Dave Piscitello:	So I'll give you an example of of a situation in the real world where people actually like to see historical data. Consider a CARFAX, okay, what you do is you consider you'd like to understand at least some aspects of the domain, maybe not the – not who owned the car, but whether or not the car, you know, was in a collision. So that's something you could get from a CARFAX.
	In the domain world, what you might want to know is whether the domain had ever appeared on a domain block list or on some URL block list. Because you may not want to instantiate a domain and go through the cost and exercise of putting up a Web site only to discover that people can't get to it because their spam filters are blocking it.
Jim Galvin:	Jim Galvin. Just a quick you said CARFAX, but for the historical record and those not from the U.S., understand, I mean, that's a service that just records the history of a vehicle through its vehicle identification number, and it tracks any as much as that stuff has been reported in the public domain, ownership and accidents and things that have happened to a vehicle. And it's a service that many used car dealers actually provide you with your car, or you can pay for it directly and get that information if you want it.
Dave Piscitello:	In fact, it's I'm glad I brought that up, because it is also a good example of how a service can have granular access controls. What you have is the VIN number, the vehicle identification number. That's what you use to do the search. There are data that identify all the previous owners in most motor vehicle databases. That data is not shared with the car dealers and it's not shared with the people who are querying CARFAX. That data is hidden from those.
	The only data that is shared is, you know, collision data. So it's an example of how you can go into a database, you can say, "This part is out of bounds, you know, and, you know, is protected by privacy rights, and this other part that actually would help you make an informed decision about buying this car is can be made publicly available."
Edmon Chung:	This is Edmon. I'm just curious. And that's a private service; right the CARFAX?
Dave Piscitello:	Yes, it's a commercial yes, commercial offering. It's a business, yes.

Edmon Chung:	Edmon again. So and also, you were just to clarify Rafik's question is that, I guess, Dave, when you mentioned that there are services out there, those are for WHOIS, now coming back to WHOIS those are also commercial services that are being provided, not something ICANN is requiring registrars
Dave Piscitello:	Absolutely. That's correct. Well, they're commercial services. And, in fact, they don't have that granularity that I just talked about. So what you are worried about or expressing concern about today, Rafik, is actually available if you are willing to pay. But a lot of things are available on the Internet, to almost any person, if you are willing to pay. But a lot of things are available on the internet, to almost any person, if you are willing to pay.
Julie Hedlund:	Do we have additional questions? Do we have any questions in the chat room?
	End of recorded material