



Internet New Zealand (Inc)

Submission to ICANN

on

Technical Evolution of the WHOIS

15 February 2011
Public Version (there is no confidential version)

For further information, please contact:

Jay Daley	.nz Registry Services	jay@nzrs.net.nz
Keith Davidson	InternetNZ	keith@internetnz.net.nz
Debbie Monahan	Domain Name Commission Ltd	dnc@dnc.org.nz

Table of Contents

1	Introduction.....	1
2	Process.....	1
3	Terminology	2
4	Lack of authentication and access control mechanism.....	3
5	Lack of query rate limiting mechanism.....	3
6	No indication of response encoding.....	4
7	Lack of standardisation in query, output and error messages.....	5
8	Incomplete support for Internationalised registration data and IDNs	5

I Introduction

- 1.1 This submission is from InternetNZ (Internet New Zealand Inc).
- 1.2 InternetNZ is a membership-based, non-partisan, not-for-profit charitable organisation responsible for the administration of the .nz top level domain.
- 1.3 Our mission is to protect and promote the Internet for New Zealand; we advocate the ongoing development of an open and uncaptureable Internet, available to all New Zealanders.
- 1.4 InternetNZ has two wholly-owned charitable subsidiaries to whom management, operation and regulation of the .nz top level domain are delegated. These are:
 - 1.1.1 .nz Registry Services, the Registry
 - 1.1.2 Domain Name Commission Limited, the Regulator

2 Process

- 2.1 This submission is in response to the request for feedback presented at the Cartagena workshop on technical evolution of the WHOIS. We have based our submission on both the document and presentation provided for that session.
- 2.2 We should note at the outset our strong concern with the process being followed by ICANN on this issue. The WHOIS is not a technical issue, it is a core service provided by many key organisations in our industry and has provided a very specific and well understood service for many years. Any change to this service should come from the position of understanding what the service is attempting to achieve and how that might be improved or amended. It should not be driven solely by a consideration of the technology that underlies the service.
- 2.3 Further we believe it is equally inappropriate for this issue to bypass the normal channels for consultation that have been negotiated with the community over many years. Our involvement came largely by accident as a consequence of attending the workshop in Cartagena and we are concerned that others who have much to contribute will not be aware of this issue due to the bypassing of the normal consultation channels.
- 2.4 Our most serious concern is that the combination of this being introduced as a technical issue, bypassing the normal channels of consultation and focusing so heavily on replacing the WHOIS protocol, brings the perception that ICANN is aiming for a pre-determined outcome and is taking steps to minimise any objections to that on the way. Such a perception, whether mistaken or not, is corrosive and should be easily avoided by following the correct process.

3 Terminology

- 3.1 Drawing on the discussion at the Cartagena workshop, we recommend using the following terminology to describe the different layers that make up WHOIS:

Layer	Meaning
WHOIS service	The service that the end user interacts with, which may be through a standalone WHOIS protocol client or a web page that queries a WHOIS protocol server and then reformats the response.
WHOIS protocol	The protocol as defined in RFC 3912.
WHOIS data representation	The format of the query made using the WHOIS protocol and the data returned in a WHOIS protocol query.
Directory services data model	The underlying model of the data that the WHOIS service is intended to provide to the end user.
Registration data	The data that is provided to the registrar and/or registry to associate with the domain name.

- 3.2 Some elements of the terminology need not be precise. For example a WHOIS response could be regarded as either a response from the WHOIS service or the WHOIS protocol or the data representation, and so is referred to a WHOIS response for short with greater precision where needed.
- 3.3 It could reasonably be argued that the WHOIS protocol could be increased in scope to define the WHOIS data representation. However, as there are many very different representations and underlying data models in common use across the various TLD sectors and the RIRs, this document identifies changes to those rather than the protocol, as such changes may be implemented without introducing protocol compatibility issues.
- 3.4 In some limited cases, standardisation of the data representation may be addressed through changes to the protocol, such as where new features are being introduced that have not yet been implemented in a variety of ways.
- 3.5 Some examples of this layer terminology applied to current practice are:
- 3.5.1 The gTLD registries (or registrars for thin registries) have converged on a single set of registration data that is represented in a single directory services data model but each registry (or registrar for thin registrars) has a different WHOIS data representation.
 - 3.5.2 Many ccTLDs have a unique set of registration data, a unique directory services data model and a unique WHOIS data representation.
 - 3.5.3 At one time, each of the RIRs accepted a different set of registration data, but used the same directory services data model and WHOIS data representation to publish them.

3.5.4 Some ccTLDs allow the end user to select a different data representation by specifying a flag in the WHOIS query, though there is no standardisation on flags.

3.6 The list of problems identified at the workshop can now be categorised by the layer(s) at which they apply:

4 Lack of authentication and access control mechanism

4.1 This issue applies to the WHOIS service and not the WHOIS protocol, as the addition of authentication and access control is a change to the service that WHOIS provides to end users. Only if such a change were agreed does it become a question of how to add this, by extending the WHOIS protocol or replacing it.

4.2 It is our view that authentication and access control is not a requirement of the WHOIS service, which is by common use and practice, over many years a public, unauthenticated service.

4.3 We have no doubt that the requirement exists for authenticated access to registration data but this should be provided through a different directory service to the WHOIS service that operates in addition to the WHOIS service rather than replacing it.

4.4 We discuss the problems that arise from adding authentication and access control to the WHOIS service in further detail below in the section on IRIS.

5 Lack of query rate limiting mechanism

5.1 This issue applies to the WHOIS service and not the WHOIS protocol, as the decision on whether to limit queries, what query rate limits to apply and how to apply those limits is a decision made by the provider of the WHOIS service.

5.2 We note that many WHOIS service providers have successfully added very different and sophisticated query rate limiting mechanisms at the TCP/IP layer without any change to the protocol. We further note that some WHOIS service providers have chosen to use the WHOIS data representation to flag their query rate limiting policy and provide feedback to the end user on their current query rates in relation to that policy.

5.3 It is possible that a requirement exists to standardise the query rate limit policy of the WHOIS service within the gTLD sector, but that is not for us to comment on, other than to note such standardisation would not require any change to the protocol and that if were implemented as a change to the protocol then that would be in direct conflict with the existing query rate limiting policies of many ccTLD.

5.4 We conclude that there is no requirement to change the WHOIS protocol to build query rate limiting into it as that has been successfully tackled by many

WHOIS service providers. It would be accurate to say that there is no lack of a query rate limiting mechanism.

6 No indication of response encoding

- 6.1 We agree that this is a notable issue with the WHOIS protocol arising from a WHOIS service requirement to support a wide variety of character sets in the data representation.
- 6.2 We note that there are a number of approaches to solving this issue, some of which do not require a change to the protocol:
 - 6.2.1 Switch the WHOIS service to another protocol that can support flagging of the encoding.
 - 6.2.2 Standardise the following three elements of encoding:
 - 6.2.2.1 the default encoding for all WHOIS queries (a change to the WHOIS protocol);
 - 6.2.2.2 the default encoding of all WHOIS responses (a change to the WHOIS protocol);
 - 6.2.2.3 the flags supplied with the query requesting different encodings (a standardisation of the data representation or a change to the WHOIS protocol)
 - 6.2.3 Add an indicator to the beginning of WHOIS responses that indicates the encoding (a standardisation of the WHOIS data representation or a change to the WHOIS protocol).
- 6.3 It is worth noting that it is often overlooked that WHOIS queries can be parameterised as the string sent in the query is not defined by the protocol. A number of registries do that already to enable choice of different character encoding:
 - 6.3.1 .dk uses "--charset=latin-1" or "--charset=utf-8"
 - 6.3.2 .no use "-c utf-8"
 - 6.3.3 .jp uses "/e" to switch from ISO-2022-JP to ASCII
- 6.4 We note that standardising the elements of encoding is the path of least resistance and may possibly be implemented with minimal impact for existing WHOIS services, though further research is required to correctly assess the impact. We note that with this mechanism the only way that the encoding of a stored WHOIS response can be determined is if that encoding is recorded as an attribute of that response, for example by storing the query that generated that response.
- 6.5 It is our view that standardising the elements of encoding or and/or adding an indicator to the beginning of WHOIS responses will solve this issue and that no replacement of the WHOIS protocol is required.

7 Lack of standardisation in query, output and error messages

- 7.1 These elements collectively make up the data representation and the decision on whether not to standardise the data representation is an issue for the WHOIS service, not the WHOIS protocol.
- 7.2 We note that different TLDs have very different data representations and have as yet, not noted or considered any requirement to standardise these other than in the limited case of RIRs.
- 7.3 We further note that the WHOIS service is, and always has been, primarily intended to be human readable above all else.
- 7.4 Some WHOIS service providers have chosen to support computer readability through the use of key-value pairs or a precisely documented data representation. This simplifies the task of systems that retrieve WHOIS data and reformat it for presentation to the end user, or use the data provided in some other way.
- 7.5 What may be less well known is that other WHOIS providers do not use key-value pairs and provide a data representation that can be presented to the end user without alteration.
- 7.6 It is possible that the requirement exists to standardise the data representation for the gTLD sector, but that is not for us to comment on. We are not aware of any such requirement in the ccTLD sector.
- 7.7 If this were to be a requirement then standardising the data representation and underlying data model is sufficiently difficult across TLDs that it is probably only realistic to expect global standardisation of some key query flags.
- 7.8 It is our view that no changes are needed to the WHOIS protocol to support this standardisation.

8 Incomplete support for Internationalised registration data and IDNs

- 8.1 This is a requirement of the WHOIS service that highlights an issue with the data model and subsequent data representation, but is limited to gTLDs as a number of ccTLDs have full support for internationalised registration data and IDNs within the scope of their policies. Any ccTLD that has not yet added such support to their WHOIS service has no external constraints that prevent this.
- 8.2 For gTLDs the most important step is to agree changes to the registration data, which could then be left to registries and registrars to build into their data model and data representations, or could be implemented in a standardised data model and data representation. We make no comment as to which of these alternative approaches should be adopted.

- 8.3 We conclude that this change does not require any change to the WHOIS protocol.

9 IRIS

- 9.1 It is our view that IRIS as a replacement for the WHOIS protocol is long dead and resurrecting it as such is pointless, but it is useful to explore why IRIS failed to get any uptake to help understand the issues better

- 9.2 IRIS failed because it made two important changes to the WHOIS service model and the community was unwilling to implement those changes as there was no consensus on the requirement for them outside of the technical community that developed the protocol.

- 9.3 The first of these changes was the addition of authentication and access control. As explained above, this was antithetical to the common understanding of the WHOIS service as a public and unauthenticated service. More seriously though, was the level of resources needed by service providers to build the backend to support such an authentication feature. Many providers decided that would require significant resources yet was of very low priority and consequently did not implement IRIS. Those that did found little or no interest from service consumers for whom the WHOIS service was sufficient and for whom to the resources required to support IRIS could not be justified.

- 9.4 The second of these changes was the use of XML to produce output that is far less human readable than most WHOIS data representations. A whole industry has built up around the current de facto standards of WHOIS data representation that relies on it being simple to process. While XML is well supported by software libraries it more complex, more verbose and less readable than most WHOIS data representation and consequently provided notable disincentives to change.

10 Additional directory services

- 10.1 We should be clear that we have no objections to considering the development of an additional directory service, based on a different protocol from WHOIS that sets out to achieve very different things from WHOIS, so long as this is not conflated with the issues around WHOIS.

- 10.2 Such a service need not put human readability and simplicity of the data model at the fore as that is catered for by the WHOIS service. Also, such a service may well have requirement for authentication and access control.

- 10.3 Some examples of potential requirements for such a service include:

- 10.3.1 Register searching as provided by a number of registries under tightly controlled circumstances.
- 10.3.2 Access to private data for law enforcement.
- 10.3.3 High volume data download.

10.3.4 Availability checker.

- 10.4 A service like this, if designed right, could see TLDs adopt the parts of it that fit with their policies in a sufficiently standardised way that the end use of the system benefits.

11 Conclusion

- 11.1 We noted at the start of this submission our concerns with the process for considering these issues.
- 11.2 We see no evidence that the WHOIS protocol needs to be replaced.
- 11.3 We see no evidence that the WHOIS protocol needs to change other than for the limited issue of indicating character encoding.
- 11.4 We recognise that there may be sufficient requirements to begin work on a new service to provide different functionality from the WHOIS service and that this should use an alternative protocol, possibly even a revised form of IRIS.

With many thanks for your consideration,

Yours sincerely,

InternetNZ