

Cross Registry Internet Service Protocol (CRISP) Requirements

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Internet registries expose administrative and operational data via varying directory services. This document defines functional requirements for the directory services of domain registries and the common base requirements for extending the use of these services for other types of Internet registries.

Table of Contents

1.	Introduction	3
1.1.	Background	3
1.2.	Requirements Scope	3
1.3.	Requirements Specification	3
2.	Internet Registry Communities	4
2.1.	Domain Name System Registries	4
2.1.1.	Domain Registries	4
2.1.2.	Domain Registrars	5
2.2.	Other Registries	5
2.2.1.	Regional Internet Registries	5
2.2.2.	Local Internet Registries	5
2.2.3.	Internet Routing Registries	5
2.2.4.	Incident Coordination Contact Registries	6
2.3.	Implementers	6
2.4.	End Users	6
2.4.1.	Internet Resource Registrants	6
2.4.2.	Service Providers and Network Operators	6
2.4.3.	Intellectual Property Holders	7
2.4.4.	Law Enforcement	7
2.4.5.	Certificate Authorities	7
2.4.6.	DNS Users	7

2.4.7.	Abusive Users	7
2.5.	Other Actors	8
3.	Functional Requirements	8
3.1.	Base Functions	8
3.1.1.	Mining Prevention	8
3.1.2.	Minimal Technical Reinvention	8
3.1.3.	Standard and Extensible Schemas	9
3.1.4.	Level of Access	9
3.1.5.	Client Processing	10
3.1.6.	Entity Referencing	10
3.1.7.	Decentralization	10
3.1.8.	Query of Access Permission	11
3.1.9.	Authentication Distribution	11
3.1.10.	Base Error Responses	11
3.1.11.	Query Distribution	12
3.1.12.	Protocol and Schema Versioning	12
3.1.13.	Relay Bag	13
3.1.14.	Privacy Labels	14
3.2.	Domain Specific Functions	14
3.2.1.	Lookups	14
3.2.2.	Searches	15
3.2.3.	Information Sets	16
3.2.4.	Serialization Support	17
3.2.5.	Result Set Limits	17
3.2.6.	DNS Delegation Referencing	17
3.2.7.	Distribution for Domain Registry Types	18
3.2.8.	Data Omission	18
3.2.9.	Internationalization	19
4.	Feature Requirements	19
4.1.	Client Authentication	19
4.2.	Referrals	20
4.3.	Common Referral Mechanism	20
4.4.	Structured Queries and Responses	20
4.5.	Existing Schema Language	20
4.6.	Defined Schemas	20
5.	Internationalization Considerations	20
6.	IANA Considerations	20
7.	Security Considerations	20
	Normative References	21
	Informative References	21
	URIs	21
A.	Glossary	23
B.	Acknowledgements	24
B.1.	Forums	24
B.2.	Working Group	24
B.3.	Contributions	25

Intellectual Property Statement	25
Author's Address	25
Full Copyright Statement	26

1. Introduction

1.1. Background

The expansion and growth of the Internet has seen the registry function of a traditionally centralized and managed Network Information Center become the responsibility of various autonomous, functionally disparate, and globally distributed Internet registries. With the broadening number of Internet registries, the uses of their administrative directory services have expanded from the original and traditional use of the whois [6] protocol to include the use of whois outside the scope of its specification, formal and informal definitions of syntax, undocumented security mechanisms, the use of other protocols, such as rwhois [5], to fulfill other needs, and proposals for the use of other technologies such as LDAP [4] and XML.

1.2. Requirements Scope

The scope of the requirements captured in this document relate to the directory services of Internet registries and their related communities (Section 2.3, Section 2.4, and Section 2.5). This scoping specifically targets the requirements of domain name registries (Section 2.1). The requirements for other registry types will be made available in other memos. The requirements are of both the current use of these directory services and the desired functionality based on input from relevant forums (Appendix B.1). These requirements are not specific to any protocol. Terms used in the definition of requirements in this document may be found in the glossary (Appendix A).

The scope of the requirements in this document are also restricted to access of data from Internet registries. Requirements for modification, addition, or provisioning of data in Internet registries are out of the scope of this document.

1.3. Requirements Specification

The requirements captured in this document are for the purpose of designing technical specifications. The words used in this document for compliance with RFC 2119 [3] do not reference or specify policy and speak only to the capabilities in the derived technology. For instance, this document may say that the protocol "MUST" support certain features. An actual service operator is always free to disable it (and then to return an error such as "permission denied".)

Requirements in this document specifying the capabilities of the protocol required for proper interaction between a client and a server will be specified with the "MUST/SHOULD" language of RFC 2119 [3]. This document also contains language relating to the interaction of a client with multiple servers to form a coherent, cross-network service. Such service requirements will not be described using RFC 2119 language.

While individual servers/service operators may not support all features that the protocol can support, they must respect the semantics of the protocol queries and responses. For example, a server should not return referrals if it does not have referent data.

2. Internet Registry Communities

The Internet registries are composed of various communities which provide scope for the requirements in this document. These communities can be generalized into the following categories: registries, registrars, implementers, end-users, and other actors.

2.1. Domain Name System Registries

2.1.1. Domain Registries

Domain registries are responsible for the registration of domains for use with DNS [1] and forward lookups (i.e., does not include the .ARPA domain). These registries have typically served two main domain functions: as the registry for a gTLD or as a registry for a ccTLD. In some instances, one entity will operate multiple TLD's, both of the gTLD and ccTLD type. A gTLD or ccTLD domain registry operator may be a governmental entity, non-governmental, non-commercial entity, or a commercial entity.

Some ccTLD's have second-level domain registrations similar in nature to gTLD's or have distinctly separate entities operating second-level domain registries similar in nature to gTLD's within the ccTLD.

Domain registries usually follow one of two models for conducting registrations of domains. The "thick" model is the more traditional model. In a "thick" domain registry, the registry contains both the operational data for the domain and the contact data (Appendix A) for the domain. In this model, the registry is typically the interface to the domain registrant but may also interface with the domain registrant through domain registrars. The "thin" model domain registry contains only operational data for domains. In the "thin" model, contact data for the domain are maintained by a domain registrar.

Domain registries not described in this section (Section 2.1.1) are not the subject of this document and may have requirements that are out of scope for this subject matter.

2.1.2. Domain Registrars

Domain registrars accept domain registrations from registrants on behalf of domain registries, both "thick" and "thin". In a "thin" model registry/registrar system, a domain registrar maintains the contact data of a domain while the registry maintains the operational data of a domain. In a "thick" model registry/registrar system, a domain registrar passes both the operational data and contact data to the registry. Domain registrars may register a domain on behalf of a registrant in more than one domain registry.

2.2. Other Registries

This section describes Internet registries other than those listed in Section 2.1. These descriptions are not definitive and this list is not absolute. They are provided in this document for informational purposes only.

2.2.1. Regional Internet Registries

Regional Internet Registries (RIR's) administer the allocation of IP address space and autonomous system numbers. Each RIR serves a specific geographic region, and collectively they service the entire Internet. Each RIR is a membership-based, non-profit organization that facilitates and implements global addressing policy based on the direction of their regional community.

2.2.2. Local Internet Registries

Local Internet Registries (LIR's) and National Internet Registries (NIR's) are sub-registries of RIR's and coordinate the same functions of the RIR's for smaller, more specific geographic regions, sovereign nations, and localities.

2.2.3. Internet Routing Registries

Internet Routing Registries are routing policy databases. Their purpose is to provide information helpful in administering Internet routers. Frequently, the syntax and contents are defined by RPSL [7].

IRR's are operated by academic, commercial, governmental, and other types of organizations, including several of the RIR's. The contents of the databases vary and reflect the needs of the users directly

served (e.g., an ISP may look up route entries, added by their customers, to decide whether to accept specific route advertisements they receive).

Unlike RIR and domain registry data, IRR data is often duplicated between separate organizations. The IRR data has the unique characteristics of being largely available through other sources (i.e., it is advertised by the Internet routing protocols) and most often having a common data format, RPSL.

2.2.4. Incident Coordination Contact Registries

Incident coordination contact registries allow operators of network resources such as network infrastructure, network names, or network services to register contact information for the purpose of providing a means of incident notification. Using this type of registry, an operator of network resources are provided information for contacting the operator of another network resource from which an incident may be occurring.

2.3. Implementers

Implementers of client software are often either affiliated with large network operators, registry operators, or commercial entities offering value-added services, or are general citizens of the Internet. Much of the client software for use with the directory services of Internet registries is either freely available, open source, or both, or available as a service. Implementers of server software are often affiliated with operators or commercial entities specializing in the out-sourcing of development for Internet registries.

2.4. End Users

This section describes the many types of end-users. Individuals and organizations may have multiple roles and may concurrently occupy many of the categories.

2.4.1. Internet Resource Registrants

Entities given authority over an Internet resource via purchase, lease, or grant from an Internet registry, either directly or via the services of a registrar.

2.4.2. Service Providers and Network Operators

Service providers and network operators provide connectivity, routing, and naming services to many other entities, some commercial

and some non-commercial, both large and small. Their operational and administrative staff often interact with Internet registries on behalf of other end-users. Service providers and network operators interact with all of the Internet registry operators outlined in this document on a frequent and consistent basis. For example, network operators use the directory services of Internet registries to determine contact information for network resources that have technical problems.

2.4.3. Intellectual Property Holders

A number of parties, such as trademark, service mark and intellectual property holders, individuals, governments and other geopolitical entities, have some legal rights on certain alphanumeric strings.

They use the directory services of Internet registries, mostly domain registries and registrars, for purposes of maintaining and defending claims to domain names consistent with applicable laws and regulations.

2.4.4. Law Enforcement

Law enforcement agencies use the directory services of Internet registries to find information used to carry out the enforcement of laws within their jurisdictions.

2.4.5. Certificate Authorities

Certificate authorities use the directory services of Internet registries as part of their verification process when issuing certificates for Internet named hosts.

2.4.6. DNS Users

Users of the Internet have client software that resolves domain names to IP addresses and IP addresses to domain names. Often when trouble occurs in the resolution process of DNS, these users trouble shoot system problems with the aid of information from the directory services of Internet registries.

2.4.7. Abusive Users

The administrative directory services of Internet registries are often the target of practices by abusive users. Using information obtained from Internet registries, abusive users undertake certain activities that are counter to the acceptable use of the information as intended by a registry, registrar, or registrant. Many times, these practices violate law in the jurisdiction of the user,

registry, registrar, or registrant. One example is the use of Internet registry information for the use of sending unsolicited bulk or commercial email.

2.5. Other Actors

Requirements must also consider the positions and policies of other actors on the use of Internet registry directory services. These actors include governments, non-governmental policy-setting bodies, and other non-governmental organizations.

3. Functional Requirements

Functional requirements describe an overall need or process for which the directory service is used by an Internet registry to fulfill its obligations to provide access to its respective customers, members, or other constituents. This section describes requirements in the manner specified in Section 1.3.

3.1. Base Functions

This section describes basic directory service protocol requirements for Internet registries. Additional requirements, specific to domain registries, are described in Domain Specific Functions (Section 3.2).

3.1.1. Mining Prevention

In order to prevent the inappropriate acquisition of data from an Internet registry's directory service, many servers will limit the amount of data that may be returned in a fixed time period from a server to a client. This will most likely be especially true for anonymous access uses (see Section 3.1.4).

The limits placed on differing types of data or applied depending upon access status will most likely differ from server to server based on policy and need. Support for varying service models in the effort to limit data and prevent data mining may or may not have a direct impact on the client-to-server protocol.

3.1.2. Minimal Technical Reinvention

The protocol **MUST NOT** employ unique technology solutions for all aspects and layers above the network and transport layers. The protocol **SHOULD** make use of existing technology standards where applicable. The protocol **MUST** employ the use of network and transport layer standards as defined by the Internet Engineering Task Force. The protocol **MUST** define one or more congestion-aware transport mechanisms for mandatory implementation.

3.1.3. Standard and Extensible Schemas

3.1.3.1. Protocol Requirement

The protocol MUST contain standard schemas for the exchange of data needed to implement the functionality in this document. In addition, there MUST be a means to allow the use of schemas not defined by the needs of this document. Both types of schemas MUST use the same schema language. The schemas MUST be able to express data elements with identifying tags for the purpose of localization of the meaning of the identifying tags.

3.1.3.2. Service Description

The client-to-server protocol must define a standard set of data structures or schemas to be used when exchanging information. It must also pose the ability to allow for the use of newer data structures that are currently not foreseen by this specification. In both cases, the description and specification of both types of data structures or schemas must be done in the same way (i.e., the same schema language).

The schemas must also be capable of "tagging" data with a unique identifier. This identifier can then be used to localize the name of that type of data. For instance, a piece of data may have the value "Bob" and its type identified with the number "5.1". Client software could use this to display "Name: Bob" in an English locale or "Nombre: Bob" in a Spanish locale.

3.1.4. Level of Access

3.1.4.1. Protocol Requirement

The protocol MUST NOT prohibit an operator from granularly assigning multiple types of access to data according to the policies of the operator. The protocol MUST provide an authentication mechanism and MUST NOT prohibit an operator from granting types of access based on authentication.

The protocol MUST provide an anonymous access mechanism that may be turned on or off based on the policy of an operator.

3.1.4.2. Service Description

Server operators will offer varying degrees of access depending on policy and need. The following are some examples:

- o users will be allowed access only to data for which they have a relationship
- o unauthenticated or anonymous access status may not yield any contact information
- o full access may be granted to a special group of authenticated users

The types of access allowed by a server will most likely vary from one operator to the next.

3.1.5. Client Processing

The protocol MUST be capable of allowing machine parsable requests and responses.

3.1.6. Entity Referencing

There MUST be a mechanism for an entity contained within a server to be referenced uniquely by an entry in another server.

3.1.7. Decentralization

3.1.7.1. Protocol Requirement

The protocol MUST NOT require the aggregation of data to a central repository, server, or entity. The protocol MUST NOT require aggregation of data indexes or hints to a central repository, server, or entity.

3.1.7.2. Service Description

Some server operators may have a need to coordinate service in a mesh or some other framework with other server operators. However, the ability to operate a CRISP compliant server must not require this.

3.1.8. Query of Access Permission

3.1.8.1. Protocol Requirement

The protocol MUST provide a mechanism allowing a client to determine if a query will be denied before the query is submitted according to the appropriate policies of the operator.

3.1.8.2. Service Description

Because usage scenarios will differ depending on both policy and type of service, some server operators may want to provide the ability for a client to predetermine its ability to retrieve data from a query. However, some operators will not allow this for security reasons, policy restrictions, or other matters.

3.1.9. Authentication Distribution

3.1.9.1. Protocol Requirement

The protocol MUST NOT require any Internet registry to participate in any authentication system. The protocol MUST NOT prohibit the participation by an Internet registry in federated, distributed authentication systems.

3.1.9.2. Service Description

Some server operators may have a need to delegate authentication to another party or participate in a system where authentication information is distributed. However, the ability to operate a CRISP compliant server must not require this.

3.1.10. Base Error Responses

The protocol MUST be capable of returning the following types of non-result or error responses to all lookups and searches:

- o permission denied - a response indicating that the search or lookup has failed due to insufficient authorization.
- o not found - the desired results do not exist.
- o insufficient resources - the search or lookup requires resources that cannot be allocated.

3.1.11. Query Distribution

3.1.11.1. Protocol Requirement

The protocol MUST NOT prohibit a server from participating in a query distribution system.

3.1.11.2. Service Description

For lookups and searches requiring distribution of queries, the client must be allowed to distribute these queries among the participants in an established mesh of server operators. It is not a requirement that the protocol enable the discovery of servers, but cooperating servers should be able to intelligently handle distribution with its established mesh. Individual server operators will respond to all queries received according to their policies for authentication, privacy, and performance.

However, the ability to operate a CRISP compliant server must not require the participation in any query distribution system.

3.1.12. Protocol and Schema Versioning

3.1.12.1. Protocol Requirements

The protocol MUST provide a means by which the end-systems can either identify or negotiate over the protocol version to be used for any query or set of queries.

All resource-specific schema MUST provide a version identifier attribute which uniquely and unambiguously identifies the version of the schema being returned in the answer set to a query.

3.1.12.2. Service Description

The service should allow end-systems using different protocol versions to fallback to a mutually supported protocol version. If this is not possible, the service must provide a meaningful error which indicates that this is the specific case.

The service must suggest negotiation and/or recovery mechanisms for clients to use when an unknown schema version is received.

3.1.13. Relay Bag

The term "bag" in this section describes a flexible container which may contain unspecified data.

3.1.13.1. Protocol Requirement

When issuing a referral, the protocol MUST be capable of supplying a relay bag from the server to the client, and the protocol MUST be capable of allowing the client to submit this relay bag with a query to the referred server. The use of the relay bag MUST be OPTIONAL. The protocol MUST NOT make any assumptions regarding the contents of the relay bag, but the relay bag MUST be described using the schema language of the protocol.

The protocol MUST provide different error messages to indicate whether the bag is of unrecognized format (permanent failure), if it contains unacceptable data (permanent failure), or if it contains data that means processing is refused at this time (transient failure).

There MUST be no more than one bag per referral. The protocol MUST NOT make an association or linkage between successive bags in a referral chain.

The client MUST pass the bag as part of any query made to a referrant server as a result of a referral.

3.1.13.2. Service Description

In some models where service coordination among participating server operators is utilized, there might be needs to allow a referring server to pass operator-to-operator coordination data along with the referral to the referent server. Such needs might be auditing or tracking. This feature requirement allows a server to pass to the client a flexible container of unspecified data ("bag") that the client should pass to the referent server. The bag has no meaning to the client.

3.1.14. Privacy Labels

3.1.14.1. Protocol Requirement

When a value in an answer to a query is given, the protocol MUST be capable of tagging the value with the following labels:

1. do not redistribute
2. special access granted

The protocol MAY define other values for this purpose, but MUST define values defined above at a minimum. The protocol MUST be capable of attaching these labels concurrently.

3.1.14.2. Service Description

Internet registries will have varying policies regarding the access to their data. Some registries may grant certain classes of users with access to data that would not normally be given to most users. In these cases, registries may want to tag the values in these entries with labels specifying the responsibilities accompanying these special user rights.

3.2. Domain Specific Functions

These functions describe requirements specifically needed by domain registries (Section 2.1.1) and domain registrars (Section 2.1.2). Requirements specific to other registries (Section 2.2) MUST be specified separately. No compliant server operator is required to support the functions required by every registry type.

3.2.1. Lookups

3.2.1.1. Protocol Requirement

The protocol MUST contain the following lookup functions:

1. Contact lookup given a unique reference to a contact of a resource.
2. Nameserver lookup given a fully-qualified host name or IP address of a nameserver.
3. Domain lookup given a fully-qualified domain name.

See Section 3.2.3 for the requirements regarding the expected return values.

3.2.1.2. Service Description

These lookups are all single index queries and should produce zero or only one entity.

Depending on the policy and need of an Internet registry, a server operator may not allow all or any of these lookups to return part or all of the information. See Section 3.2.3.

3.2.2. Searches

3.2.2.1. Protocol Requirement

The protocol MUST contain the following search functions:

1. Domain name search given an exact match or reasonable subset of a name. This search SHOULD allow for parameters and qualifiers designed to allow better matching of internationalized domain names and SHOULD allow for both exact and partial matching within the limits of internationalized domain names. This search SHOULD NOT require special transformations of internationalized domain names to accommodate this search. This search MUST provide a means to narrow the search by names delegated under a particular TLD.
2. Domain registrant search by either exact name or partial name match with the ability to narrow the search to registrants of a particular TLD.
3. Domains hosted by a nameserver given the fully-qualified host name or IP address of a nameserver.

See Section 3.2.3 for the requirements regarding the expected return values.

3.2.2.2. Service Description

Depending on the policy and need of an Internet registry, a server operator may not allow all or any of these searches to return part or all of the information. See Section 3.1.4. Access to information resulting from these searches may also be limited, depending on policy, by quantity. Section 3.2.5 describes these types of restrictions.

Some Internet registries may also be participating in a query distribution system. See Section 3.1.11.

3.2.3. Information Sets

3.2.3.1. Protocol Requirements

The data sets for contacts, nameservers, and domains MUST be able to express and represent the attributes and allowable values of registration requests in domain registration and provisioning protocols.

The schema MUST be capable of expressing the following information for domains:

- o activation status
- o registrant
- o nameservers
- o technical, billing or other contacts
- o registry delegating the domain
- o registrar for the domain

The data set for domains MUST be able to express arbitrary textual information for extensions on an individual operator basis. Examples of such information are license agreements, authorized use policies, extended status notifications, marketing/for sale notices, and URI references to other sources.

3.2.3.2. Service Description

It is not expected that every Internet registry supply all of the information spelled out above, however the schemas employed by the protocol must be capable of expressing this information should a registry need to provide it.

The following sections describe requirements relative to the use of schemas with respect to individual registry need and policy:

- o Section 3.2.8
- o Section 3.2.5
- o Section 3.1.4
- o Section 3.1.1

3.2.4. Serialization Support

The schemas used by the protocol SHOULD be capable of off-line serialization

Off-line serialization allows for implementation independent operations such as backup and recovery, load-balancing, etc. This MAY also make possible, in whole or in part, data escrow capabilities and other usages, however such usages are out of the scope of this document.

3.2.5. Result Set Limits

3.2.5.1. Protocol Requirement

The protocol MUST contain a feature, used at the discretion of a server operator, to allow a server to express to a client a limit on the number of results from searches and lookups. When returning result sets, the protocol MUST be able to make the following distinctions:

1. an empty result set.
2. a result set truncated for the purpose of improving performance bottlenecks.
3. a result set truncated to comply with Section 3.1.1

3.2.5.2. Service Description

Client software will operate more usefully if it can understand reasons for the truncation of result sets. Of course, some Internet registries may not be able to expose their policies for the limiting of result sets, but, when it is possible, clients will have a better operational view. This may eliminate re-queries and other repeated actions that are not desirable.

3.2.6. DNS Delegation Referencing

3.2.6.1. Protocol Requirement

The protocol MUST use the delegation authority model available in DNS [1] as the primary means for determining the authoritative source for information regarding domains or any other objects when applicable.

3.2.6.2. Service Description

The intent of this requirement is to have clients use the DNS delegation model to find servers authoritative for resources instead of using a master or central server containing pointer information. In other words, when a resource is naturally mapped by DNS, the desired behavior is to consult the DNS to find an authoritative server containing information about that resource. Using 'example.com', the authoritative server for information about example.com according to the registrant of that domain may be found by querying the DNS zone for example.com. To find the registry information for example.com, the DNS zone for .com should be queried.

There are cases where resources will not naturally map into the DNS delegation hierarchy. This requirement is not meant to force such a mapping.

3.2.7. Distribution for Domain Registry Types

3.2.7.1. Protocol Requirement

The protocol MUST NOT prohibit the distribution of data to exclude any of the registry/registrar models stated in Section 2.1.1. The protocol MUST be capable of expressing referrals and entity references between the various models described in Section 2.1.1.

3.2.7.2. Service Description

Depending on the domain registry/registrar model in use, technical data for a domain may only reside in one server while contact data for the same domain may only reside in a server operated by a separate entity. However, in many uses, this is not the situation. Therefore, the service must accommodate for the various registration distribution models of domain registry types described in Section 2.1.1 while complying with Section 3.1.7.

3.2.8. Data Omission

3.2.8.1. Protocol Requirement

When a value in an answer to a query cannot be given due to policy constraints, the protocol MUST be capable of expressing the value in one of three ways:

1. complete omission of the value without explanation
2. an indication that the value cannot be given due to insufficient authorization

3. an indication that the value cannot be given due to privacy constraints regardless of authorization status

The protocol MAY define other values for this purpose, but MUST define values defined above at a minimum.

3.2.8.2. Service Description

Internet registries will have varying constraints regarding their ability to expose certain types of data, usually social information. Server operators must have the ability to accommodate this need while client software will be more useful when provided with proper explanations. Therefore, depending on policy, a server operator has a choice between not returning the data at all, signaling a permission error, or indicating a privacy constraint.

3.2.9. Internationalization

The schema defining domain related resources MUST conform to RFC 2277 [2] regarding textual data. In particular, the schema MUST be able to indicate the charset and language in use with unstructured textual data.

The protocol MUST be able to support multiple representations of contact data, with these representations complying with the requirements in Section 3.2.3. The protocol MUST be able to provide contact data in UTF-8 and SHOULD be able to provide contact data in US-ASCII, other character sets, and capable of specifying the language of the data.

4. Feature Requirements

Feature requirements describe the perceived need derived from the functional requirements for specific technical criteria of the directory service. This section describes requirements in the manner specified in Section 1.3.

4.1. Client Authentication

Entities accessing the service (users) MUST be provided a mechanism for passing credentials to a server for the purpose of authentication. The protocol MUST provide a mechanism capable of employing many authentication types and capable of extension for future authentication types.

4.2. Referrals

To distribute queries for search continuations and to issue entity references, the protocol MUST provide a referral mechanism.

4.3. Common Referral Mechanism

To distribute queries for search continuations and to issue entity references, the protocol MUST define a common referral scheme and syntax.

4.4. Structured Queries and Responses

To provide for machine consumption as well as human consumption, the protocol MUST employ structured queries and responses.

4.5. Existing Schema Language

To provide structured queries and responses and allow for minimal technological reinvention, the protocol MUST employ a pre-existing schema language.

4.6. Defined Schemas

To provide for machine consumption as well as human consumption, the protocol MUST define schemas for use by the structured queries and responses.

5. Internationalization Considerations

Requirements defined in this document MUST consider the best practices spelled out in [2].

6. IANA Considerations

IANA consideration for any service meeting these requirements will depend upon the technologies chosen and MUST be specified by any document describing such a service.

7. Security Considerations

This document contains requirements for the validation of authenticated entities and the access of authenticated entities compared with the access of non-authenticated entities. This document does not define the mechanism for validation of authenticated entities. Requirements defined in this document MUST allow for the implementation of this mechanism according best common practices.

The requirement in Section 3.1.4 must be weighed against other requirements specifying search or lookup capabilities.

This document contains requirements for referrals and entity references. Client implementations based on these requirements SHOULD take proper care in the safe-guarding of credential information when resolving referrals or entity references according to best common practices.

This document contains requirements for the distribution of queries among a mesh of participating service providers. Protocols proposed to meet these requirements must be able to protect against the use of that distribution system as a vector of distributed denial of service attacks or unauthorized data mining.

Normative References

- [1] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [2] Alvestrand, H., "IETF Policy on Character Sets and Languages", BCP 18, RFC 2277, January 1998.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Informative References

- [4] Wahl, M., Howes, T. and S. Kille, "Lightweight Directory Access Protocol (v3)", RFC 2251, December 1997.
- [5] Williamson, S., Kusters, M., Blacka, D., Singh, J. and K. Zeilstra, "Referral Whois (RWhois) Protocol V1.5", RFC 2167, June 1997.
- [6] Harrenstien, K., Stahl, M. and E. Feinler, "NICNAME/WHOIS", RFC 954, October 1985.
- [7] Alaettinoglu, C., Villamizar, C., Gerich, E., Kessens, D., Meyer, D., Bates, T., Karrenberg, D. and M. Terpstra, "Routing Policy Specification Language (RPSL)", RFC 2622, June 1999.

URIs

- [8] <<http://www.ietf.org/proceedings/00dec/00dec-41.htm>>
- [9] <<http://www.ietf.org/proceedings/01aug/51-40.htm>>

- [10] <<http://www.uwho.verisignlabs.com/Final-WhoIsPanel-Aug15-Resume.pdf>>
- [11] <http://www.ripe.net/ripe/meetings/archive/ripe-40/minutes/min_database.html>
- [12] <<http://www.nanog.org/mtg-0110/lookup.html>>

Appendix A. Glossary

- o TLD: Initials for "top level domain." Refers to domains in DNS [1] that are hierarchically at the level just beneath the root.
- o ccTLD: Initials for "country code top level domain." TLD's which use one of the two character country codes defined by ISO.
- o gTLD: Initials for "generic top level domain." TLD's that do not use one of the two character country codes defined by ISO.
- o contact data: Data containing names and contact information (i.e., postal addresses, phone numbers, e-mail addresses) of humans or legal entities.
- o operational data: Data necessary to the operation of networks and network related services and items.
- o RIR: Initials for "regional Internet registry."
- o IRR: Initials for "Internet routing registry."
- o forward lookup: a DNS lookup where a domain name is resolved to an IP address.
- o reverse lookup: a DNS lookup where an IP address is resolved to a domain name.
- o mining: In the context of this document, this term is specific to data mining. This is a methodical process to obtain the contents of directory service, usually as much as possible, not relevant to any immediate need. Data mining is often not a practice welcomed by registry operators.

Appendix B. Acknowledgements

B.1. Forums

The proceedings of the following public forums were used as input to the scope and requirements for this document:

- o whois BOF of the 49th IETF [8]; December 10-15, 2000; San Diego, CA, USA
- o whoisfix BOF of the 51st IETF [9]; August 5-10, 2001; London, England
- o First UWho Consultation [10]; August 15, 2001; Washington, DC, USA
- o Second UWho Consultation; November 15, 2001; Marina del Rey, CA, USA
- o Third UWho Consultation; November 19, 2001; Washington, DC, USA
- o DNR WG of RIPE 40, October 1-5, 2001; Prague, Czech Republic
- o Database WG of RIPE 40 [11]; October 1-5, 2001; Prague, Czech Republic
- o General Session of NANOG 23 [12]; October 21-23; Oakland, CA, USA
- o DNR WG of RIPE 41, January 14-18, 2002; Amsterdam, The Netherlands
- o Database WG of RIPE 41, January 14-18, 2002; Amsterdam, The Netherlands
- o NANOG 24 Universal Whois BOF, February 10-12, 2002; Miami, Florida
- o CENTR General Assembly, February 21-22, 2002; Rambouillet, France
- o CRISP BOF of the 53rd IETF, March 17-22, 2002, Minneapolis, Minnesota, USA

B.2. Working Group

This document is a work item of the Cross-Registry Internet Service Protocol (CRISP) Working Group in the Applications Area of the IETF. Discussions for this working group are held on the email list ietf-not43@lists.verisignlabs.com. To subscribe to this email list, send email to ietf-not43-request@lists.verisignlabs.com with a subject line of "subscribe". Archives of this list may be found out <http://lists.verisignlabs.com/pipermail/ietf-not43/>.

B.3. Contributions

Comments, suggestions, and feedback of significant substance have been provided by Leslie Daigle, Mark Kosters, Ted Hardie, Shane Kerr, Cathy Murphy, Stephane Bortzmeyer, Rick Wesson, Jaap Akkerhuis, Eric Hall, Patrick Mevzek, Marcos Sanz, Vittorio Bertola, George Michaelson, and Tim Christensen.

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Author's Address

Andrew L. Newton
VeriSign, Inc.
21355 Ridgetop Circle
Sterling, VA 20166
USA

Phone: +1 703 948 3382
EMail: anewton@verisignlabs.com; anewton@ecotroph.net

Full Copyright Statement

Copyright (C) The Internet Society (2004). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.