DNS Extensions
Internet-Draft
Updates: 2782 (if approved)
Intended status: Standards Track
Expires: April 25, 2011

O. Gudmundsson
Shinkuro Inc.
A. Hoenes
TR-Sys

October 22, 2010

DNS SRV Usage Guidelines and Owner Name Recommendations
draft-gudmundsson-dnsext-srv-clarify-02

Abstract

The DNS SRV record has been specified in RFC 2052 and RFC 2782 for use in dynamic service discovery for a domain. These two RFCs did not clearly specify an IANA registry for the names of the services and their underlying protocols. This document clarifies RFC 2782 regarding the formation and use of the Service Prefix in the owner name of SRV records, based on the unified IANA "Service Name and Transport Protocol Port Number" registry and provides a few other corrections, clarifications, and usage recommendations.

Table of Contents

1. Introduction

1.1. Synopsis of RFC 2782

The SRV resource record (RR) was originally introduced in RFC 2052 [RFC2052] as an experimental extension to the Domain Name System (DNS). It proved a highly valuable addition for service location and provisioning. The SRV record was thus standardized in RFC 2782 [RFC2782]. The main difference between these two versions is the use of the underscore ("_") prefix character in names to avoid naming conflicts between service names and regular names.

The presentation form of the SRV resource record (RR type 33), including the recommended structure of owner names for its use, is as follows [RFC2782]:

```
_Service._Proto.Name SRV Priority Weight Port Target
```

The owner name is formed from the fully qualified domain name (FQDN) (indicated by "Name") of the domain that is offering the service by prepending a "Service Prefix" that consists of two more DNS labels. The first label prepended to "Name" (indicated by "_Proto") denotes the protocol over which the application protocol is carried; it is henceforth designated as the "Protocol Label".
The final (least significant, i.e. leftmost) label prepended (indicated by "_Service") denotes the Service/Application that is being offered; it is henceforth designated as the "Service Label".

The RDATA portion of the SRV resource record is comprised of the 16-bit unsigned integer components 'Priority', 'Weight', and 'Port', and a domain name, 'Target':
The Priority and Weight fields are used for selecting among multiple SRV records at the same owner name (see RFC 2782 [RFC2782]).
The Port field is the port number over which the service is provided at the Target name, which is an FQDN of the node hosting the service.

1.2. Problem Statement

RFC 2782 says that the source of names for "Service" and "Proto" is "Assigned Numbers" (STD2) or a locally defined repository.

However, upon reflection, both alternatives do not seem to make particular sense, and therefore this has mislead authors of specifications and implementers and users of service discovery:

- o The STD2 series of documents was obsoleted by RFC 3232 [RFC3232] and IANA registration publication was handed over to on-line registries maintained by IANA. Unfortunately it is not explicitly

explained in RFC 2782 which section of STD2 it is referring to, nor does RFC 3232 help. By common knowledge, RFC 2782 referred to the Keyword columns of the "Protocol Numbers" and "Port Numbers" IANA registries, respectively.

The corresponding sections in the last edition of STD 2 [RFC1700] were "Assigned Internet Protocol Numbers" (pp. 8-10), "Well Known Port Numbers" (pp. 16-39) and "Registered Port Numbers" (pp. 39-55).

But IANA also maintains the "Protocol and Service Names" registry (with file name "service-names"), dedicated to the "Official Protocol Names as these appear in the Domain Name System WKS records and the {legacy} NIC Host Table", and pointing to RFC 952 [RFC0952] for the description of use. (Note that WKS records contain a protocol-specific bit map of "Well Known Services" offered at a host. The protocol and service names in that registry are used in the presentation form of WKS records and translated by DNS servers to protocol numbers and bit positions for the well-known ports of the services; the reverse translation is also applied in DNS software where WKS RR presentation format is generated. The WKS record format worked pretty well when well-known ports were limited to the range up to 255, but long after the publication of the DNS Standards, that range has been expanded by one of the predecessors of RFC 1700, RFC 1340 [RFC1340] to up to 1023, which made WKS records much less practical -- see [Clnup].) Since this registry unabashedly mixes names for protocols and services and should only contain entries for applications that have been assigned a well-known (system) port number, it should be clear that this registry could not have been in mind at the time RFC 2782 has been written. However, anecdotal evidence has shown that this wasn't actually clear to all parties all over the years.

As SRV records contain the port number where each server provides the service, the outmost utility of SRV RRs is for services that do not have a registered port number. (Services offered at a port number different from the registered default port number or multiple instances of a service offered at different port numbers are the next most important use cases.) Further, the number of ports available is small compared to the possible number of service names that could be registered, and so the set of services that can make use of SRV based service discovery is potentially much larger than the set of services with assigned port numbers. Therefore, the "Port Number" registry needed a more strict registration policy and is not the proper place for registering service names for use with SRV resource records.

o On the other hand, having locally defined lists of service and/or protocol names (as suggested by the text of RFC 2782) would equally allow to list the full service information in such local databases and thus make the usage of SRV records redundant. In any case, this scenario is not applicable for publicly available services where potential clients are not under the control of the authority offering the services, and hence most probably would have no access to the proper "locally provided" information. The reader is reminded that locally maintained database solutions generally scale very poorly, and that this once was the major motive for the deployment of the Domain Name System. Therefore, locally defined lists are not suitable sources for the names appearing in the Service Prefix of SRV owner names -- at least for services offered for larger than local scope.

1.3. Objective for this Document and its Companions

In an effort to streamline and fix mistakes in assigning service names and port numbers over the years, an effort has been undertaken to create a unified "Service Name and Transport Protocol Port Numbers" registry [I-D.ietf-tsvwg-iana-ports] where all service registrations will take place in the future.

[[RFC-Editor: Please replace (throughout the entire document) 'yyyy' by the RFC number assigned to Ref. [I-D.ietf-tsvwg-iana-ports] wherever "RFC yyyy" is used as a shorthand for that document.]]

This document updates RFC 2782 [RFC2782] in clarifying how to use the registrations in that registry and other elements to construct standard SRV owner names.
Furthermore, this document proposes an extension method for specific applications that need a finer granularity in specifying protocol stacks underlying the application protocol or have other important reasons to qualify the service instance provided.
Additionally, usage considerations for SRV based service lookup and the provisioning of SRV records are presented, based on experience gained since the publication of RFC 2782.

A companion document of this one, "Service Registration Fixes" [Clnup], collects from various RFCs and IANA documents registrations that ought be in the "Service Name and Transport Protocol Port Number" registry, and it specifies appropriate actions in moving the registrations over to the new registry. It also gives advice for authors wishing to update existing specifications for particular applications that had produced SRV usage guidelines in a manner that is no more admissible under the unified rules of RFC yyyy [I-D.ietf-tsvwg-iana-ports] and this document. These recommendations include the option to make use of the extension scheme suggested in the Appendix.

Note: RFC 5507 [RFC5507] discusses the use of "underscore labels" in the DNS more generally, from the architectural point of view of the IAB. Most such uses are not related to SRV records and hence will not conflict with the SRV Protocol Label; thus, there will not be namespace conflicts.

1.4. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

For clarity and precision, this document uses the following capitalized terms:

Service Name
: The identifier for a particular service registered in the IANA "Service Name and Transport Protocol Port Number" registry [I-D.ietf-tsvwg-iana-ports]. See Section 4 for details.

  Note: Traditionally, this term has been used only for the symbolic names assigned to "Well Known Ports" (aka "Assigned Port" or "System Ports") in the IANA "Protocol and Service Names" registry in support of the "NIC Hosts Table" (RFC 810 [RFC0810] and its successor, RFC 952 [RFC0952]) and the DNS WKS resource record [RFC1035]; this has been done so in all "Assigned Numbers" RFCs since RFC 900 ([RFC0900], [RFC0923], [RFC0943], [RFC0960], [RFC0990], [RFC1010], [RFC1060], [RFC1340], [RFC1700]); RFC 1060 has used this term in the sense we use it now, but restricted to the "Unix Standard services"; later RFCs have used the term "keyword" and "short name" for this purpose; [I-D.ietf-tsvwg-iana-ports] re-establishes and expands the general use of this term.

Service Domain
: The FQDN of the domain for which a particular instance of a service is provided (or sought by a prospective client).

Service Prefix (or: SRV Service Prefix)
: The service-specific relative domain name (additional labels) prepended to the Service Domain to form the owner name of the SRV resource record(s) pointing to the service instance(s). As usual, in the DNS presentation format (e.g., in a DNS zone file) all domain name labels are concatenated using the period (".") as a separator character.

In RFC 2782, the Service prefix consists of exactly two labels (described right off below), and this specification does not change this. The labels contained in a Service Prefix adhere to the common naming hierarchy rules for the DNS, specifying the most significant part (the protocol immediately carried over IP) on the righthand side and the particular application/ service as the least significant part on the lefthand side.

Protocol Label (or: SRV Protocol Label)
The most significant (righthand) label of the Service Prefix. It indicates the transport protocol to which the 'Port' element of the SRV RDATA serves as the demultiplexer value.
See Section 3 for details.

Service Label (or: SRV Service Label)
The more specific, second label of the Service Prefix. It indicates the particular service by incorporating the Service Name registered for this service/application.
See Section 4 and Appendix A for details.

2. General Considerations for SRV Service Prefixes

SRV Service Prefixes SHOULD consist of exactly two labels, as specified in RFC 2782 [RFC2782].

NOTE: There are cases where (in practice or under consideration) additional labels are prepended. The extension scheme in Appendix A is offered as a guideline for an alternative to such multi-level hierachical Service Labels that conforms to the general rules. But since this extension proposal is currently non-normative, we cannot reasonably ban present use with a "MUST" in place of the "SHOULD" above.

If the specification of a particular application/service indicates that this service is to be provided at a specific registered port number and does not mention DNS SRV based service discovery, prospective clients SHOULD NOT assume the existence of SRV records for this application, unless the client has hints available that indicate otherwise -- for instance by means of configuration. In this case, only the rules of Sections 3 and 4 apply.

If the specification of a particular application/service does mention dynamic service discovery based on DNS SRV records but does not specify otherwise in a precise and unambiguous manner, the rules of Sections 3 and 4 apply. In this case, a prospective client SHOULD look up the DNS for the appropriate SRV records for the intended Service Domain before trying to contact the service at its assigned port number (if any).

Some services can be carried equivalently in different encapsulations using higher-level "substrate" protocols like HTTP, BEEP, SOAP, SIP, or XMPP, some of which in turn can be carried over different transport protocols. In this case, it is possible that certain servers for an application only support specific protocol stacks, or that a Service Domain provides a different set of servers for each protocol stack. This is likely to happen during the introduction of new transport schemes for an application. When generic middleware layers are used, there commonly is no practical method to let such plug-in middleware layers chose among them based on packet content, or to perform a negotiation of the particular substrate protocol. Consequently, there occasionally is a need to specify such details in the SRV Service Prefix. Since the Protocol Label is intentionally restricted (Section 3), this information has to be carried in the Service Label for this application.

Within the semantics of the tree-oriented naming scheme used in the DNS, there are two possibilities to achieve this goal for a particular application:

a. If there is only a small number of "substrate" protocols to distinguish, it is RECOMMENDED that the application designers register multiple Service Names with IANA in the "Service Name and Transport Protocol Port Number" registry [I-D.ietf-tsvwg-iana-ports], which usually will start with the same characters and contain a suffix attached with a hyphen embedded in the name.

    One disadvantage of this solution is the length limitation imposed on the Service Name by RFC yyyy [I-D.ietf-tsvwg-iana-ports] (15 characters, see Section 4). Also, RFC yyyy recommends that -- notwithstanding some very popular, and mostly legacy cases -- each service ought to be represented by a *single* Service Name. If the service is the same and only the substrate used is different, this solution arguably is in conflict with RFC yyyy.

b. If the number of substrate stacks to be supported is larger, or if the 'canonical' labels that the application designers prefer would result in exceeding the maximum length of Service Names, or if multi-layer substrates shall be represented, or if the "unique Service Name per service" argument is deemed more important, an extension scheme is needed for the construction of longer, more-specific Service Labels for such exceptional services.

The non-normative Appendix A proposes a scheme for extended Service Labels that does not pollute the namespace of Service Names and hence adheres to the uniqueness and collision-freeness requirements of RFC yyyy [I-D.ietf-tsvwg-iana-ports]. The basic Service Name can be easily extracted from these extended Service Labels, as they employ a separator character not allowed in Service Names to attach qualifier(s) to a 'basic' Service Label as specified in Section 4.

Application designers who want to make use of that scheme MUST independently and unambiguously specify the application of that scheme for their service and ensure that a reference to the document containing this specification (if different from the document specifying the basic service) is added to the entry for that service in the IANA "Service Name and Transport Protocol Port Number" registry, following the registration procedures of [I-D.ietf-tsvwg-iana-ports].

To re-iterate, absent such explicit, registered specification, the rules of Section 4 (and, of course, Section 3 as well) still apply in these cases.

Application protocol designers ought to keep service discovery simple. The fewer alternatives a prospective client has to consider and the fewer choices it has to find suitable SRV records, the faster the service discovery can be performed, because it needs fewer DNS queries on average and thus incurs less latency. If alternative transports and/or variations of the service are really needed to be distinguished by service discovery, to foster interoperablity there SHOULD always be defined a default version supported by all clients and servers, and configured in the DNS, thus allowing an orderly fallback in case client preferences cannot be accommodated.

3. Protocol Labels

SRV Protocol Labels MUST be formed by prepending an underscore character ("_") to the name of an IETF transport protocol that

- is registered in the IANA "Protocols" registry [RFC5237];

- makes use of 16-bit unsigned integer demultiplexing values, at least at the server side of its transport associations -- these are traditionally denoted as "[server] port numbers";

- is supported by the IANA "Service Name and Transport Protocol Port Number" registry as specified in [I-D.ietf-tsvwg-iana-ports] or by future IESG-endorsed amendments to that registry.

Like all domain labels, SRV Protocol Labels are matched in a case-insensitive manner.

Currently, the following SRV Protocol Labels are defined :

Protocol	SRV Protocol Label	References
TCP	_tcp	[RFC0793] [RFC4614]
UDP	_udp	[RFC0768]
UDP-lite	_udp	[RFC3828]
DCCP	_dccp	[RFC4340]
SCTP	_sctp	[RFC4960]

Table 1

Note that, although they have been assigned different protocol numbers in the IANA Protocol Numbers registry [RFC5237]), the IANA "Service Name and Transport Protocol Port Number" registry [I-D.ietf-tsvwg-iana-ports] does not make a distinction between 'classical' UDP [RFC0793] and UDP-lite [RFC3828], and hence both protocol variations are represented by the same SRV Protocol Label. In practice, this makes it impossible to define and support SRV-based service discovery for applications/services that can be carried over either UDP or UDP-lite, since both variants would have to be reperesented by SRV records at the same node in the DNS database, with no means to distinguish between the different transport protocols inside the SRV RRs. Arguably, this raises another obstacle to the deployment and use of UDP-lite and makes a smoth migration from UDP to UDP-lite almost impossible.

4. Standard Service Labels

Absent a normative document specifying otherwise, the SRV Service Label for a specific service/application SHOULD be in the form of a "Standard SRV Service Label" as specified in this section.

Standard SRV Service Labels MUST be formed by prepending an underscore character ("_") to the Service Name of an application protocol registered in the IANA "Service Name and Transport Protocol Port Number" registry [I-D.ietf-tsvwg-iana-ports].

Note that in order to make use of DNS-based service discovery using SRV records, the application/service does *not* need an IANA-assigned port number (which would also be filed in the IANA "Service Name and Transport Protocol Port Number" registry).

The port numbers carried in the RDATA portion of the SRV records are locally assigned within the Service Domain; in case the service indeed has an IANA-assigned (default) port number, the locally assigned port number MAY coincide with that default port number, unless the documentation of the service specifies otherwise.

Like all domain name labels, SRV Service Labels are matched in a case-insensitive manner.

[I-D.ietf-tsvwg-iana-ports] restricts IANA-registered Service Names to 15 characters in "ldh-syntax", informally:

- all characters MUST be either letters from A to Z, decimal digits from 0 to 9, or hyphens ("-");

- there MUST be at least one letter in a Service Name;

- the first and the last character MUST be a letter or digit, not a hyphen;

- two hyphens MUST NOT be adjacent in a Service Name.

More formally, this can be restated in ABNF [RFC5234]:

```
serv-label     = std-serv-label

std-serv-label = USC service-name

service-name   = 1*15(L-D-H)
                   ; also conforming to the <token> rule below

token          = *(1*DIGIT [HYPHEN]) ALPHA *([HYPHEN] ALPHANUM)

L-D-H          = ALPHA / DIGIT / HYPHEN
ALPHANUM       = ALPHA / DIGIT   ; A-Z, a-z, 0-9
HYPHEN         = %x2d            ; "-" (hyphen)
USC            = %x5f            ; "_" (underscore)
ALPHA          = <See RFC 5234>  ; A-Z, a-z
DIGIT          = <See RFC 5234>  ; 0-9
```

Note: In case the <service-name> rule turns out to be in conflict with RFC yyyy [I-D.ietf-tsvwg-iana-ports], the latter has precedence.

However, in the -07 draft version of [I-D.ietf-tsvwg-iana-ports], the formal syntax definition (ABNF) does not agree with the prose; for clarity, a consistent version of the ABNF is presented above, which is expected to be in the next version of that I-D as well.

The length restriction is there to support use of Service Names in other contexts than SRV Service Labels (which, as all DNS labels, can be 63 octets long). The other restrictions assure that Service Names cannot be confused with numbers (e.g., port numbers) or port range notation (like nn-mmm) and that they follow the common restrictions for domain labels that prohibit conflicts with domain name internationalization efforts.

Appendix A offers a scheme for "Extended SRV Service Labels" that can be adopted by service specifications that cannot contend with these restrictions, and which seek for a versatile naming scheme not violating the provisions of [I-D.ietf-tsvwg-iana-ports].

5. Service Discovery Client Considerations

Implementations making use of dynamic service discovery based on DNS SRV records for a particular application/service will construct a prioritized list of applicable Service Prefixes, following the guidelines in the two subsections below.

To form the QNAME(s) for DNS SRV lookup, each Service Prefix is concatenated (with a period as the label separator character) to the FQDN of the Service Domain the application is interested in.

Depending on application strategy and perhaps local policy (configuration), the DNS queries with QCLASS=IN and QTYPE=SRV can be performed serially or in parallel -- to decrease the latency in the case higher priority queries do not succeed in finding matching SRV record(s).

The answers obtained are processed as specified in RFC 2782 [RFC2782], subject to the preferences of the service client for transport and/or "substrate" protocols. If necessary, the Target domain names obtained are then queried for address records (i.e., at the time of this writing, A and/or AAAA RRs) to determine the network layer addresses to be contacted over the corresponding transport protocol using the port number contained in the 'Port' element of the respective SRV record.

Update to RFC 2782:

RFC 2782 [RFC2782] on page 7 imprecisely indicates the client "[SHOULD] ... try to connect to the (protocol, address, service)" tuple. Since at the transport layer the port number needs to be used, not the service (name?), and to let the order of the tuple components follow the protocol layers, it should refer to "(address, protocol, port)" tuples.

5.1. Protocol Label Selection

Some services are defined to operate over only one transport protocol. In this case the application MUST use the appropriate transport label in forming the Service Prefix for SRV record lookup.

If a service can operate over multiple transport protocols, then the specification of the service might indicate an order of preference, or local policy may supply it. Any service can provide a specific order in the Notes column of the IANA "Service Name and Transport Protocol Port Number" registry entry, during registration or via registration update [I-D.ietf-tsvwg-iana-ports].

Absent such information, the following priority order SHOULD be used as appropriate, based on which transports are defined and registered for the service, and which are supported by the client: _udp, _tcp, _sctp, _dccp.

5.2. Service Label Selection

Absent service specific documentation saying otherwise or hints available to the client (e.g., by configuration), the following recommendations SHOULD be followed.

If an application/service has a single registered Service Name, a prospective client uses the Standard Service Label derived from it according to Section 4.

Contrary to the strong recommendation in RFC 2782 [RFC2782], several legacy services have been assigned more than one Service Name in the past. For instance, the well-know "systat" service is also referred to as "users" service, the DNS service, "domain", is also referred to as "nameserver" or "dns", "rlp" is also denoted as "resource", the WhoIs service is identified by "whois" and "nicname", the web service uses "http" or "www", "kerberos" is also indicated by "kdc", and the PCMail service has been assigned the equivalent Service Names "pcmail-srv" and "pcmail".

In all these cases, the "Service Name and Transport Protocol Port Number" registry will clearly indicate (in the Notes column of the registry entries) which name is the primary one and which names are considered aliases, including pointers from the entries for alias Service Names to the entry of the primary Service Name.

Unless a prospective client has specific hints available (e.g., by configuration) indicating that a specific alias name ought to be tried preferentially, the primary Service Name SHOULD be used, and it also SHOULD be tried if the lookup of an alias name fails.

If a service/application supports different well marked instances identified by different Service Names or a related specification has introduced the usage of extended Service Labels for discovery of this service (e.g., making use of the scheme specified in Appendix A), the specifications SHOULD always also define a single default service instance, and hence a default Service Label or Service Name (that can be used to construct the corresponding Standard Service Label). Prospective clients MAY follow any service instance selection policy desired (by implementation, deployment, or configuration), but SHOULD be prepared to fall back to the default service instance if the SRV record lookup for preferred service instance(s) fails.

6. Provisioning of SRV records

DNS zone administrators SHOULD support (and encourage) the provisioning of SRV records related to the basic domains they manage. Dynamic DNS Update ([RFC2136], [RFC3007]) is an option, but this is out of scope for this document.

The structure of SRV owner names makes it possible (but not necessary) to delegate the per-protocol subdomains (given by the defined Protocol Labels) related to a Service Domain. If delegated, because various services are offered over different protocols, preferably all the subdomains SHOULD be served within the same administrative domain (or even on the same authoritiaive servers), thereby simplifying administrative procedures for the maintenance of the SRV records.

For services with alias names available as well, DNS zone managers SHOULD strongly encourage the use of the primary Service Names for the formation of the SRV Service Labels. Applications registering SRV names in the DNS SHOULD always use the primay Service Name. If Service Prefixes based on "nicknames" are wanted, for consistency and maintenance reasons it is RECOMMENDED that CNAME records be placed at the corresponding owner names instead of SRV records, thus redirecting DNS lookups to a single SRV RRset bound to the primary Service Name.

[RFC2782] goes at considerable length to explain why SRV RRs are not suitable for short-term load balancing purposes, and that therefore SRV records can have reasonably long TTL values.

The combination of 'Priority' and 'Weight' in SRV RRs has proven ardous to make proper use of, and inconsistent client behavior makes it unlikely that hypothetical strategies for fine-tuned load distribution can be achieved in practice by populating a two-dimensional variety of (Priority, Weight) tuples in any SRV RRset.

Experience has shown that fixing a 'standard' value for 'Priority' and varying only the 'Weight' values is much more likely to influence client behavior in a deterministic way.

Note: The SRV 'Target' is an FQDN, not an IP address. In the past, there have been attempts to encourage the use of IP addresses for 'Target', but that is not possible. Some implementations have been observed to accept an IPv4 address for 'Target' and store it as a 4-label domain name, causing the trailing (top-level) label to become numeric, which is not allowed.

7. Security Considerations

This document does not have any specific security implications. However it is hoped that the more orderly, and more frequent use of SRV-based dynamic service discovery, based on the rules clarified in this documents and the establishment of a unified service registry, will provide valuable information for administrators and security policy makers, to the benefit of the overall security of the Internet.

For the general security considerations of SRV resource records, see RFC 2782 [RFC2782]. The use of DNSSEC [RFC4033] to digitally sign zones publishing SRV records provides data integrity to the DNS lookup and prevents spoofing.

All security considerations in RFC yyyy [I-D.ietf-tsvwg-iana-ports] related to port numbers apply to applications/protocols that use SRV based service discovery rather than assigned (server) port numbers.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

[I-D.ietf-tsvwg-iana-ports]
Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", draft-ietf-tsvwg-iana-ports-07 (work in progress), October 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.

[RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

9.2. Informative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[RFC0810] Feinler, E., Harrenstien, K., Su, Z., and V. White, "DoD Internet host table specification", RFC 810, March 1982.

[RFC0900] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 900, June 1984.

[RFC0923] Reynolds, J. and J. Postel, "Assigned numbers", RFC 923, October 1984.

[RFC0943] Reynolds, J. and J. Postel, "Assigned numbers", RFC 943, April 1985.

[RFC0952] Harrenstien, K., Stahl, M., and E. Feinler, "DoD Internet host table specification", RFC 952, October 1985.

[RFC0960] Reynolds, J. and J. Postel, "Assigned numbers", RFC 960, December 1985.

[RFC0990] Reynolds, J. and J. Postel, "Assigned numbers", RFC 990, November 1986.

[RFC1010] Reynolds, J. and J. Postel, "Assigned numbers", RFC 1010, May 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC1060] Reynolds, J. and J. Postel, "Assigned numbers", RFC 1060, March 1990.

[RFC1340] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1340, July 1992.

[RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, October 1994.

[RFC2052] Gulbrandsen, A. and P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2052, October 1996.

[RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.

[RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.

[RFC3232] Reynolds, J., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, January 2002.

[RFC3828] Larzon, L-A., Degermark, M., Pink, S., Jonsson, L-E., and G. Fairhurst, "The Lightweight User Datagram Protocol (UDP-Lite)", RFC 3828, July 2004.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, March 2006.

[RFC4614] Duke, M., Braden, R., Eddy, W., and E. Blanton, "A Roadmap for Transmission Control Protocol (TCP) Specification Documents", RFC 4614, September 2006.

[RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

[RFC5237] Arkko, J. and S. Bradner, "IANA Allocation Guidelines for the Protocol Field", BCP 37, RFC 5237, February 2008.

[RFC5507] IAB, Faltstrom, P., Austein, R., and P. Koch, "Design Choices When Expanding the DNS", RFC 5507, April 2009.

Appendix A. Extended SRV Service Labels

This non-normative Section defines a versatile extension scheme for SRV Service Labels that can be incorporated by reference in

specifications of service discovery procedures for applications that cannot contend with Standard Service Labels as specified in Section 4. This way, the specification below becomes normative for these (and only these) particular applications/services.

A.1. Motivation and Solution Space

Sections 2 and 4 identify scenarios of applications/services that might reasonably wish to use an extended scheme for forming their Service Labels, in particular to identify protocol layers (encapsulation layers, "substrate" protocols) stacked between the transport protocol and the application protocol itself, without having to register a multitude of Service Names in the "Service Names and Port Numbers" registry and hence being bound to the 15-character name size limit. Other scenarios with similar requirements have been mentioned in recent work-in-progress in the IETF.

For being useful, a scheme for extended Service Labels must allow (a) to easily determine that a label obeys to this scheme and (b) to unambiguously extract the underlying Service Name and the added extension components (henceforth: Qualifiers) from the full label, whereas at the same time the scheme needs to ensure that it does not pollute the name space of Service Names and that thereby the uniqueness of registered Service Names is not disturbed.

As restated in Section 4, the syntax of Service Names from [I-D.ietf-tsvwg-iana-ports] does not allow the underscore ("_") character, which in turn already is used as a prefix character for SRV Service and Protocol Labels, where it serves to distinguish these labels from 'ordinary' domain name components. Therefore, a manifest method to construct extended Service Labels is to concatenate the given Service Name and Qualifiers, prepending each component by an underscore character.

A.2. Specification

Thus, we arrive at the following syntax for Extended Service Labels, extending the ABNF from Section 4:

```
serv-label     =/ ext-serv-label

ext-serv-label = USC service-name 1*( USC serv-qualifier )
                 ; total size limited to 63 characters

serv-qualifier = 1*(l-d-h)
                 ; also conforming to the <token> rule
```

Hypothetical example:
Assume an application with registered Service Name "foohoo" for transport over SCTP is served for the domain "ext.example". Also assume that this service needs to be qualified by the keywords (service qualifiers) (1) "barbar" and (2) "aanne".
Then the corresponding Extended SRV Service Label will be:

```
_fooho_barbar_aanne
```

and the full SRV owner name will be:

```
_fooho_barbar_aanne._sctp.ext.example
```

Application designers who want to make use of this scheme need to precisely document the <serv-qualifier> values supported and refer normatively to this section. Any such specification SHOULD indicate a mandatory-to-implement default form of the service that will be represented by the Standard Service Label for this service, per Section 4. This allows for an easy fallback strategy for clients of such service that are not interested in particular variants of the service, or when the service variant preferred by the client is not offered at a particular Service Domain and hence not represented by an SRV record in the DNS.
Absent such application-specific documentation, always the Standard Service Labels specified in Section 4 are used.

A.3. Applicability

Each application/service making use of this mechanism inherits from its registered Service Name a distinct namespace, and its designers must manage this 'private' namespace of valid Extended Service Labels according to their needs. For instance, there is no central IANA registry for such namespaces. The application is still identified by its Service Name, and the related leading standard Service Label part is to be used for policy decisions.

If the service qualifiers are used to indicate intermediate layers, the application-specific service discovery specification SHOULD specify that the qualifiers be given in protocol stacking order; if substrate protocols used have their own registered Service Name, it is strongly RECOMMENDED that these Service Names be used to identify the corresponding qualifiers.

The companion document [Clnup] contains examples of legacy use cases specified in the IETF that preferably should migrate to this scheme and thereby, it instantiates these recommendations.

Authors' Addresses

Olafur Gudmundsson
Shinkuro Inc.
4922 Fairmont Avenue, Suite 250
Bethesda, MD 20814
USA

Email: ogud@ogud.com

Alfred Hoenes
TR-Sys
Gerlinger Str. 12
Ditzingen D-71254
Germany

Email: ah@TR-Sys.de