Network Working Group	D. Crocker
Internet-Draft	Brandenburg InternetWorking
Intended status: Best Current Practice	March 30, 2011
Expires: October 01, 2011	

DNS Scoped Data Through Attribute Leaves
draft-crocker-dns-attrleaf-04

Abstract

Historically, any DNS RR may occur for any domain name. Recent additions have defined DNS leaf nodes that contain a reserved node name, beginning with an underscore. The underscore construct is used to define a semantic scope for DNS records associated with the parent domain. This note explores the nature of this DNS usage and defines the "underscore names" registry with IANA.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

The core DNS technical specifications assign no semantics to domain names or their parts, and no constraints upon which resource records (RRs) may be associated with particular names. Over time, some leaf node names, such as "www" and "ftp" have come to imply support for particular services, but this is a matter of operational convention, rather than defined protocol semantics. This freedom in the basic technology has permitted a wide range of administrative and semantic policies to be used -- in parallel -- with the DNS. Data semantics have been limited to the specification of particular resource records, on the expectation that new ones would be added as needed.
Some recent service enhancements have defined define a restricted scope for the occurrence of particular resource records. That scope is a leaf node, within which the uses of specific resource records can be formally defined and constrained. This leaf has a distinguished naming convention: It uses a reserved DNS node name that begins with an underscore. Because host names are not allowed to use the underscore character, this distinguishes the name from all legal host name. Effectively, this convention creates a space for attributes that are associated with the parent domain, one level up.
An established example is the SRV record [RFC2782] which generalizes concepts long-used for email routing by the MX record [RFC0974] [RFC2821]. The use of special DNS names has significant benefits and detriments. Some of these are explored in [RFC5507].

[Comment]: The terms "resolution context" and "scoping rules" have been suggested, in place of "semantic scope". In order to avoid concern for matters of semantics, this specification uses the term

"scoping rules", to create a focus on the mechanics being defined, rather than nuances of interpretation for the mechanism.

The scoping feature is particularly useful when generalized resource records are used -- notably TXT and SRV. It provides efficient separation of one use of them from another. Absent this separation, an undifferentiated mass of these RRs are returned to the client which then must parse through the internals of the records in the hope of finding ones that are relevant. With underscore-based scoping, only the relevant RRs are returns.
This specification discusses this enhancement, provides an explicit definition of it, and establishes an IANA registry for the reserved names that begin with underscore.

Discussion Venue: Discussion about this draft is directed to the dnsop@lists.uoregon.edumailing list of the IETF DNSOP Working Group.

2. Scaling Benefits and TXT and SRV Resource Records

Some resource records have a generic form, and support a variety of uses. Each additional use defines its own rules and, possibly, its own internal syntax and node-naming conventions to distinguish among particular types. The TXT and SRV records are the notable concern for this. Some of these approaches scale poorly, particularly when the same RR can be present in the same leaf node, but with different uses. An increasingly-popular approach, with excellent scaling properties, uses an underscore-based name to a define place in the DNS that is constrained to particular uses for particular RRs. This means that a direct lookup produces only the desired records, at no greater cost than a typical lookup.
In the case of TXT records, use for different scoping rules has developed organically and largely without coordination. One side-effect of this is no consistently distinguishable internal syntax for the records; even internal inspection might not be a reliable means of distinguishing among them. Underscore-based names therefore provide an administrative way of separating TXT records that might have different uses, but otherwise would have no syntactic markers for distinguishing among them.
In the case of the SRV RR this method of distinguishing among uses was part of the design. [RFC2782] In reality, the SRV specification defines an RR that may only be used for specific applications when there is an additional specification. So the SRV specification is best thought of as a template for future specifications. The template definition includes reference to tables of names from which underscore-names should be drawn. So, the set of <service> names is defined in terms of other IANA tables, namely any table with symbolic names. The other SRV naming field is <proto>, although its pool of names is not explicitly defined.

3. Underscore DNS Registry Function

This specification defines a registry for DNS nodes names, used to specify scope of use for specific resource records (RR). That is, a given names defines a specific, constrained context for the use of such records. This does not constrain the use of other resource records that are not specified. The purpose of the registries is to avoid collisions resulting from the use of the same underscore name, for different applications.
Structurally, the registry is defined as a single, flat table of names that begin with underscore. In some cases, such as for SRV, an underscore names might have further constraints, such as being valid only "under" some other underscore name. Semantically, this is a hierarchical model, thereby making a flat registry unexpected.
The registry requires such hierarchies to be registered as a combinatorial case analysis set, with each entry being a full sequence of underscore names. Given a naming scheme that is actually structured, this flat design is inelegant. However it has the benefit of being extremely simple, with the added advantage of being easier for readers to understand, as long as these cases are small and few.

NAME
_protoA
_service1._protoB
_service2._protoC
_service2._protoC
_service3._protoD._useX
_protoE._region._authority

Example of Underscore Names

The reasons for choosing a simplified registry design are:

- the belief that listing multi-level schemes as complete combinations will be simpler than formulating sub-tables, simples, and

- the view that requiring readers to parse through a possible hierarchy of multiple registries -- one per level -- will encourage errors.

4. DNS Underscore Registry Definition

A registry entry MUST contain:

- **Name:** Specifies a textual name for a scoped portion of the DNS. The name will usually be taken from the specification cited in the "Purpose" column and is intended for use in discussions about the entry.

DNS Label(s):
Specifies a sequence of one or more underscore names that define a single name reservation.

Constraints: Specifies any restrictions on use of the name.

RR(s): Lists the RRs that are defined for use within this scope.

References Lists specifications that define the records and their use under this Name.

Purpose: Specifies the particular purpose/use for specific RR(s), defined for use within the scope of the registered underscore name.

5. IANA Considerations

IANA is requested to establish a DNS Underscore Name Registry, for DNS node names that begin with the underscore character (_) and have been specified in any published RFC, or are documented by a specification published by another standards organization. The contents of each entry are defined in Section 3.

NAME	DNS LABEL	CONSTRAINTS	RR(s)	REFERENCES	PURPOSE
SIP TCP	_sip._tcp		NAPTR	[RFC3263]	Locating SIP Servers
SIPS TCP	_sips._tcp		NAPTR	[RFC3263]	Locating SIP Servers
SIP UDP	_sip._udp		SRV	[RFC3263]	Locating SIP servers.
SPF	_spf		TXT	[RFC4408]	
DKIM	_domainkey		TXT	[RFC4871]	Public key for verifying DKIM signature.
ADSP	_adsp._domainkey		TXT	[RFC5617]	
PKI LDAP	_PKIXREP._ldap		SRV	[RFC4386]	LDAP PKI Repository
PKI HTTP	_PKIXREP._http		SRV	[RFC4386]	HTTP PKI Repository
PKI OCSP	_PKIXREP._ocsp		SRV	[RFC4386]	OCSP PKI Repository

DNS Underscore SCOPE Name Registry (with initial values)

6. Security Considerations

This memo raises no security issues.

7. References

7.1. References -- Normative

[RFC2782]	Gulbrandsen, A., Vixie, P. and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.
[RFC3263]	Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, June 2002.
[RFC4408]	Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", RFC 4408, April 2006.

7.2. References -- Informative

[RFC2821]	Klensin, J., "Simple Mail Transfer Protocol", RFC 2821, April 2001.
[RFC0974]	Partridge, C., "Mail routing and the domain system", RFC 974, January 1986.
[RFC4871]	Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J. and M. Thomas, "DomainKeys Identified Mail (DKIM) Signatures", RFC 4871, May 2007.
[RFC5507]	Faltstrom, P. and R. Austein, "", RFC 5507, April 2009.
[RFC4386]	Boeyen, S. and P. Hallam-Baker, "Internet X.509 Public Key Infrastructure: Repository Locator Service", February 2006.
[RFC5617]	, , , , , and , "DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP)", August 2009.

Appendix A. Acknowledgements

Thanks go to Bill Fenner, Tony Hansen, Peter Koch, Olaf Kolkman, and Andrew Sullivan for diligent review of the earlier drafts.

Author's Address

Dave Crocker Crocker Brandenburg InternetWorking 675 Spruce Dr. Sunnyvale, CA 94086 USA Phone: +1.408.246.8253 EMail: dcrocker@bbiw.net URI: http://bbiw.net/