INTERNET-DRAFT
draft-ietf-dnsop-serverid-00.txt

David Conrad
Nominum, Inc.

May, 2002

Identifying an Authoritative Name Server

Status of this Memo

Abstract

A standardized mechanism to determine the identity of a name server responding to a particular query would be useful, particularly as a diagnostic aid. This document describes an identification convention used in one widely deployed implementation of the DNS protocol and proposes a slight modification to that convention aimed at addressing some implementation concerns.

1. Introduction

Determining the identity of the name server responding to a query has become more complex due primarily to the proliferation of various load balancing techniques. This document describes a convention used by one particular DNS server implementation to provide identifying information and proposes a slight modification to that convention to address concerns regarding implementation neutrality.

Note that this document makes no value judgements as to whether or not the convention in current use is good or bad; it merely documents

the covention's existence and proposes a slight redefinition of the convention to address non-technical implementation concerns.

2. Rationale

Identifying which name server is responding to queries is often useful, particularly in attempting to diagnose name server difficulties. However, relying on the IP address of the name server has become more problematic due the deployment of various load balancing solutions, including the use of shared unicast addresses as documented in [RFC3258].

An unfortunate side effect of these load balancing solutions is that traditional methods of determining which server is responding can be unreliable. Specifically, non-DNS methods such as ICMP ping, TCP connections, or non-DNS UDP packets (e.g., as generated by tools such as "traceroute"), etc., can end up going to a different server than that which receives the DNS queries.

This proposal makes the assumption that an identification mechanism that relies on the DNS protocol is more likely to be successful (although not guaranteed) in going to the same machine as a "normal" DNS query.

3. Historical Conventions

Recent versions of the commonly deployed Berkeley Internet Name Domain implementation of the DNS protocol suite from the Internet Software Consortium [BIND] support a way of identifying a particular server via the use of a standard, if somewhat unusual, DNS query. Specifically, a query to a late model BIND server for a TXT resource record in class 3 (CHAOS) for the domain name "HOSTNAME.BIND." will return a string that can be configured by the name server administrator to provide a unique identifier for the responding server (defaulting to the value of a gethostname() call). This mechanism, which is an extension of the BIND convention of using CHAOS class TXT RR queries to sub-domains of the "BIND." domain for version information, has been copied by several name server vendors.

For reference, the other well-known name used by recent versions of BIND within the CHAOS class "BIND." domain is "VERSION.BIND." A query for a TXT RR for this name will return an administratively re-definable string which defaults to the version of the server responding.

4. An Implementation Neutral Convention

The previously described use of the CHAOS class "BIND." domain has

rightly been viewed by many implementors as not being standardized nor being implementation neutral. As such, a standard mechanism to identify a particular machine among a shared unicast set of machines serving the same DNS data does not currently exist.

Since a name server conforming to [RFC1034] and [RFC1035] should support the CHAOS class and the use of TXT resource record queries in the CHAOS class to derive information about a name server has been used in several independent name server implementations, the quickest way of supporting the identification of a particular name server out of a set of name servers all sharing the same unicast prefix would likely be to standardize on the BIND convention, albeit with a slight modification to address implementation neutrality concerns.

The convention proposed here simply redefines the top level CHAOS domain to be "SERVER." instead of "BIND.". Since using the actual hostname may be considered an information leakage security risk, the use of the actual hostname of the server is discouraged and instead a unique per-server identifier should be used. As the BIND convention of "HOSTNAME" implies the use of a hostname, the domain name "ID.SERVER" is proposed. That is, a TXT RR query for "ID.SERVER." in the CHAOS class will return an administratively defined string that can be used to differentiate among multiple servers.

To make this convention useful, DNS operators wishing to identify their servers MUST put a unique string for the RDATA of the TXT record associated with the "ID.SERVER." domain in class CHAOS. Implementors MUST provide a way to disable returning identifying information. Implementors SHOULD provide a way to limit who can query for the identifying information.

The use of other names in the CHAOS class "SERVER." domain are beyond the scope of this document.

IANA Considerations

The "SERVER." domain in the CHAOS class should be reserved by IANA and a registry should be created that reserves the "ID" name. In the future, requests may be submitted for other sub-domains of "SERVER.", e.g., "VERSION.SERVER." and the IANA should take appropriate action.

Security Considerations

Providing identifying information as to which server is responding can be seen as information leakage and thus a security risk. It may be appropriate to restrict who can query for the "ID.SERVER." domain. Filtering on source address would be one way in which restrictions can be applied.

The identifer returned via an "ID.SERVER." query SHOULD NOT contain the hostname or other information that could be considered sensitive.

Acknowledgements

The technique for host identification documented here derive from practices implemented by Paul Vixie of the Internet Software Consortium in the Berkeley Internet Name Domain package. Useful comments on earlier drafts were provided by Bob Halley, Brian Wellington, Andreas Gustafsson, Ted Hardie, Chris Yarnell, and members of the ICANN Root Server System Advisory Council.

References

[RFC1034] Mockapetris, P., "Domain Names - Concepts and Facilities", RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain Names - Implementation and Specifications", RFC 1035, November 1987.

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

[RFC3258] Hardie, T., "Distributing Authoritative Name Servers via Shared Unicast Addresses", RFC 3258, April, 2002.

Author's Address

David Conrad
Nominum, Inc.
2385 Bay Road
Redwood City, CA 94063
USA

Phone: +1 650 381 6003
Fax: +1 650 381 6055
Email: david.conrad@nominum.com

Full Copyright Statement