Workgroup: DNSOP Working Group
Internet-Draft:
draft-grubto-dnsop-dns-out-of-protocol-
signalling-03
Published: 10 July 2023
Intended Status: Standards Track
Expires: 11 January 2024
Authors: C. Almond P. van Dijk M.W. Groeneweg
ISC PowerDNS SIDN
S.W.J. Ubbink D. Salzman W. Toorop
SIDN CZ.NIC NLnet Labs

DNS Out Of Protocol Signalling

Abstract

This document seeks to specify a method for DNS servers to signal programs outside of the server software, and which are not necessarily involved with the DNS protocol, about conditions that can arise within the server. These signals can be used to invoke actions in areas that help provide the DNS service, such as routing.

Currently this document serves as a requirements document to come to a signalling mechanism that will suit the use cases best. Part of that effort is to assemble a list of conditions with potential associated out of DNS protocol actions, as well as inventory and assess existing signalling mechanisms for suitability.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 11 January 2024.

Table of Contents

1. Introduction

Operators of DNS servers can benefit from automatically taking action upon certain conditions in the name server software. Some conditions can be monitored from outside the server software, but for adequate and immediate action, the server software itself should

signal about the condition immediately when it occurs to invoke action by a listener for these signals.

An example of such a condition is when all zones, from a set served from an anycasted prefix, are loaded and ready to be served. An associated action may be to start announcing a prefix route from the point-of-presence where the name server is running and to withdraw the prefix route if one of the zones cannot be served anymore. This way queries for zones will only reach the point-of-presence if the name server software can answer those queries.

Another example condition may be if an recursive resolver served from an anycasted prefix, is started and ready to serve, with the same associated action of only announcing the anycasted prefix when the recursive resolver can serve queries.

All anycasted DNS services can benefit from the mechanism alone, by the increased adequacy and reduced resources of not having to poll for a server's state. DNS services with diverse implementations will benefit from standardizing of the name server signalling.

Before coming to a specification for the mechanism, this document will serve to inventorise the already available standardized and non-standardized signalling channels and assess them for usability for out of protocol signalling.

2. Terminology and Definitions

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Conditions to be signalled

This section served to collect a list of conditions for which actions outside of the DNS protocol may be interesting. A signal will be sent if the condition is met, and also when the condition is no longer met. Some conditions take configuration parameters influencing when the conditions are met. Some conditions may contain arguments when signalled. When applicable, the parameters and arguments are given with each condition.

Some conditions may be identified from outside of the DNS server by polling for the condition. This is more resource intensive that listening for a signal, but may also be more robust. When this is the case, how the condition can be identified is provided with the condition.

3.1. The DNS server is running and can respond to queries

How to identify:

* check if the DNS server is running by doing a query to see if it responds

Action:

* Start announcing the prefix on which this zone is served with BGP A announcement may be withdrawn when the condition is no longer met.

3.2. Shutting down

How to identify:

* Maintenance, before shutting down the name server, initiate at least the BGP withdrawl

Action:

* Stop the BGP announcement of the prefix

3.3. The nameserver has crashed

How to identify:

* The name server is no longer running (or does not respond to queries, although that might also be the case when it is under an attack)

Action:

- Stop the BGP announcement of the prefix

This condition maybe only detected from outside of the DNS server.

3.4. A zone is loaded and ready to serve

How to identify:

* Query the zone to see if it responds

Argument:

* The zone that was loaded

Action:

- Start announcing the prefix on which these zones are served with BGP. A announcement may be withdrawn when the condition is no longer met.

Some name servers, when configured to notify targets when a zone is updated [RFC1996], will also notify those targets when a zone is just loaded. The notify itself may be considered an appropriate signal, although it will not be emitted when the zone is no longer served.

3.5. All zones are loaded and ready to serve

Action:

- Start announcing the prefix on which these zones are served with BGP. A announcement may be withdrawn when the condition is no longer met.

This condition may be derived from one or more "A zone is loaded and ready to serve" (Section 3.4) signals when a list of all zones served is available.

3.6. A zone is updated to a new version

How to identify:

- Query the zone's SOA record, register value and then compare to expected version

Argument:

- The zone that was updated

Action:

- Verify the zone content. Is it DNSSEC valid, does the ZONEMD validate.

Name servers can usually already signal this with NOTIFY [RFC1996]

3.7. A zone is (about to) expire

Parameter:

- The period before expiration. A value of 0 will emit the signal the moment the zone expires.

Argument:

- *The zone that is (about to) expire

Action:

- *Stop the BGP announcement of the prefix on which the zone is served. It may be reannounced when the zone becomes available again (See Section 3.4).

3.8. DNSSEC signatures are (about to) expire

Parameter:

- *The period before expiration. A value of 0 will emit the signal the DNSSEC signature expires.

Argument:

- *The zone that contains the signature

- *The resource record set owner name and type with the signature that will soon expire

Action:

- *Stop the BGP announcement of the prefix on which the zone is served. It may be reannounced when the zone becomes DNSSEC valid again.

3.9. Query rate is exceeding a threshold

Parameter:

- *The number of queries per second threshold.

Action:

- *Lengthen the AS path for the BGP announcement for a prefix, to demotivate the anycast node that receives all the queries.

- *Or if the query rate is indicating a denial of service attack, keep the BGP AS path short, to absorb the attack.

- *Signal to Security Information and Event Management SIEM and logging that problem has been observed.

3.10. Query rate increase is exceeding a threshold

Parameter:

- *The number of queries per second increase per second threshold.

Action:

- *The same actions as for "Query rate is exceeding a threshold" (Section 3.9) apply.

3.11. Extended DNS Error conditions

Parameter:

- *The Extended DNS Error conditions for which to signal [RFC8914]

Argument:

- *The Extended DNS Error condition that occurred.

Action:

- *Dependent on the DNS Error condition

4. Requirements for signalling mechanisms and channels

- *All conditions are sensitive information and should be stay either in the administered domain (for example on the local machine that is under control of the operator), or needs to be authenticated.

5. Existing signalling mechanisms and channels

What follows is a list of existing signalling mechanisms and a comparison of those channels in Section 5.5.

5.1. Notify

DNS NOTIFY [RFC1996] is an existing ubiquitous mechanism to signal zones. It is intended to target name servers, but tooling exists to listen for NOTIFY messages and trigger execution of a command when a zone is updated (See [nsnotifyd]).

Advantages:

- *Native signalling for zone updates present right now (See Section 3.6)

- *Indirect support for zone loaded (See Section 3.4)

Disadvantages:

* One available Open Source Software which lacks authentication support and is therefore only suitable for local usage

* Only two conditions are signalled.

* Does not signal when the conditions are no longer met.

5.2. D-Bus as publication channel

D-Bus is a mechanism for exchanging messages between processes local on the same machine (See [D-Bus]). The D-BUS protocol is a one-to-one protocol, but distribution of messages (or signals) to multiple other applications is carried out by a program intended for this purpose: the D-Bus *message bus*.

Advantages:

* Implementation already exists (See [Knot-DNS-3.1.6])

* Good Open Source Software library support [TODO references]

Disadvantages:

* Server needs to be started before clients making it less robust.

* Is only communicated locally to the machine

5.3. DDoS Open Threat Signaling

DDoS Open Threat Signaling (DOTS) [RFC9132] is a set of protocols for real-time signaling of threat-mitigation requests within and between different operational domains.

Advantages:

* Publish / Subscribe mechanism

* Inter-operator communications

* Authenticated

* Open Source server software exists [TODO reference go-dots]

Disadvantages:

* No Open Source client library exists? We need to get information during the upcoming hackathon at the IETF117. Current DOTS builds upon CoAP [RFC7252] for which many client library implementations exist.

5.4. MQTT

MQTT (see [MQTT-OASIS-Standard-v5]) is a lightweight publish-subscribe network protocol for messages.

Advantages:

*Network Publish / Subscribe mechanism

*Supports authentication

Disadvantages:

*Need to gain experience at the IETF117 hackathon

5.5. Observations and comparison

Method	NOTIFY	D-Bus	DOTS	MQTT
Local to machine	+	++	+	+
inter-machine	+	-	+	+
inter-operator	+	-	++	-
Publish Subscribe	-	-	++	++
Authentication	+-	-	+	+
Client library availability	NA	++	?	++

Table 1

6. Security and Privacy Considerations

Signalling MUST be performed in an authenticated and private manner.

7. Implementation Status

*Knot DNS has support for D-Bus notifications (See Section 5.2) for significant server and zone events with the "dbus-event" configuration parameter since version 3.1.6 [Knot-DNS-3.1.6]

*NSD has a feature branch [NSD-oops-branch] where work is being done on the implementation

8. IANA Considerations

This document has no IANA actions

9. Acknowledgements

We would like to thank the people of the port53 hackathon in Rotterdam for their contributions. Mainly Doris Hauser, Lars-Johan Liman, Vilhelm Prytz and Henrik Kramselund

10. Normative References

[RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, DOI 10.17487/RFC1996, August 1996, <https://www.rfc-editor.org/info/rfc1996>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <https://www.rfc-editor.org/info/rfc7252>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8914] Kumari, W., Hunt, E., Arends, R., Hardaker, W., and D. Lawrence, "Extended DNS Errors", RFC 8914, DOI 10.17487/RFC8914, October 2020, <https://www.rfc-editor.org/info/rfc8914>.

[RFC9132] Boucadair, M., Ed., Shallow, J., and T. Reddy.K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", RFC 9132, DOI 10.17487/RFC9132, September 2021, <https://www.rfc-editor.org/info/rfc9132>.

11. Informative References

[D-Bus] Pennington, H., Carlsson, A., Larsson, A., Herzberg, S., McVittie, S., and D. Zeuthen, "D-Bus Specification", February 2023, <https://dbus.freedesktop.org/doc/dbus-specification.html>.

[Knot-DNS-3.1.6] CZ.NIC, "Knot DNS - Version 3.1.6", February 2022, <https://www.knot-dns.cz/2022-02-08-version-316.html>.

[MQTT-OASIS-Standard-v5] Banks, A., Briggs, E., Borgendale, K., and R. Gupta, "OASIS Standard MQTT Version 5.0", 19 March 2019, <https://docs.oasis-open.org/mqtt/mqtt/v5.0/os/mqtt-v5.0-os.html>.

[NSD-oops-branch] NLnet Labs, "NSD feature/oops branch", May 2023, <https://github.com/NLnetLabs/nsd/tree/features/oops>.

[nsnotifyd] Finch, T., "nsnotifyd: scripted DNS NOTIFY handler", January 2022, <https://dotat.at/prog/nsnotifyd/>.

Appendix A. Implementation Status

Note to the RFC Editor: please remove this entire appendix before publication.

Knot currently uses [D-Bus] for this.

Appendix B. Change History

Note to the RFC Editor: please remove this entire appendix before publication.

*draft-grubto-dnsop-dns-out-of-protocol-signalling-03

> Rename "name server" into "DNS server" when it also applies to recursive resolvers
>
> Make a single list of conditions with per condition indicated the parameters (how they can be influenced by configuration), the arguments (the signal payload) and "how to identify" if the condition can be identified from outside of the DNS server.
>
> Removing DNS Error reporting monitoring agent as a channel to evaluate
>
> Add DOTS and MQTT as a potential signal channels for our conditions

*draft-grubto-dnsop-dns-out-of-protocol-signalling-02

> Updates after discussion during the port53 hackathon in Rotterdam.

*draft-grubto-dnsop-dns-out-of-protocol-signalling-00

> Initial version

Authors' Addresses

Cathy Almond
Internet Systems Consortium, Inc.
PO Box 360
Newmarket, NH 03857
United States of America

Phone: +1 650 423 1300
Email: cathya@isc.org

Peter van Dijk
PowerDNS
Den Haag
Netherlands

Email: peter.van.dijk@powerdns.com

Marc Groeneweg
Stichting Internet Domeinregistratie Nederland
Postbus 5022
6802EA Arnhem
Netherlands

Email: marc.groeneweg@sidn.nl

Stefan Ubbink
Stichting Internet Domeinregistratie Nederland
Postbus 5022
6802EA Arnhem
Netherlands

Email: stefan.ubbink@sidn.nl

Daniel Salzman
CZ.NIC
Czechia

Email: daniel.salzman@nic.cz

Willem Toorop
NLnet Labs
Science Park 400
1098 XH Amsterdam
Netherlands

Email: willem@nlnetlabs.nl