Network Working Group
Internet-Draft
Intended status: Informational
Expires: 28 October 2022

M. Boucadair
Orange
T. Reddy.K
Akamai
W. Pan
Huawei Technologies

26 April 2022

Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS) draft-ietf-dots-multihoming-13

Abstract

This document discusses multi-homing considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS). The goal is to provide some guidance for DOTS clients and client-domain DOTS gateways when multihomed.

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 28 October 2022.

Table of Contents

1. Introduction

In many deployments, it may not be possible for a network to determine the cause of a distributed Denial-of-Service (DoS) attack [RFC4732]. Rather, the network may just realize that some resources appear to be under attack. To help with such situations, the IETF has specified the DDoS Open Threat Signaling (DOTS) architecture [RFC8811], where a DOTS client can inform an upstream DOTS server that its network is under a potential attack and that appropriate mitigation actions are required. The DOTS protocols can be used to coordinate real-time mitigation efforts which can evolve as the attacks mutate, thereby reducing the impact of an attack and leading

to more efficient responsive actions. [RFC8903] identifies a set of scenarios for DOTS; most of these scenarios involve a Customer Premises Equipment (CPE).

The high-level base DOTS architecture is illustrated in Figure 1 ([RFC8811]):

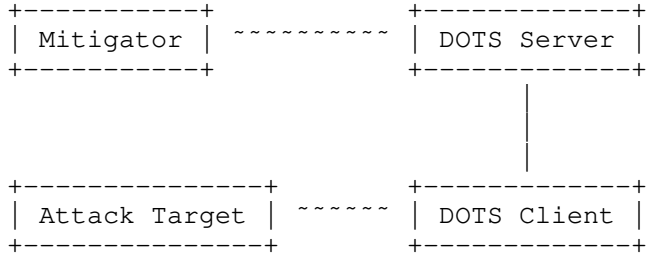

Figure 1: Basic DOTS Architecture

[RFC8811] specifies that the DOTS client may be provided with a list of DOTS servers; each of these servers is associated with one or more IP addresses. These addresses may or may not be of the same address family. The DOTS client establishes one or more DOTS sessions by connecting to the provided DOTS server(s) addresses (e.g., by using [RFC8973]).

DOTS may be deployed within networks that are connected to one single upstream provider. DOTS can also be enabled within networks that are multi-homed. The reader may refer to [RFC3582] for an overview of multi-homing goals and motivations. This document discusses DOTS multi-homing considerations. Specifically, the document aims to:

1. Complete the base DOTS architecture with multi-homing specifics. Those specifics need to be taken into account because:

   * Sending a DOTS mitigation request to an arbitrary DOTS server will not necessarily help in mitigating a DDoS attack.

   * Randomly replicating all DOTS mitigation requests among all available DOTS servers is suboptimal.

   * Sequentially contacting DOTS servers may increase the delay before a mitigation plan is enforced.

2. Identify DOTS deployment schemes in a multi-homing context, where DOTS services can be offered by all or a subset of upstream providers.

3. Provide guidelines and recommendations for placing DOTS requests in multi-homed networks, e.g.,:

   * Select the appropriate DOTS server(s).

   * Identify cases where anycast is not recommended for DOTS.

This document adopts the following methodology:

* Identify and extract viable deployment candidates from [RFC8903].

* Augment the description with multi-homing technicalities, e.g.,

  - One vs. multiple upstream network providers

  - One vs. multiple interconnect routers

  - Provider-Independent (PI) vs. Provider-Aggregatable (PA) IP addresses

* Describe the recommended behavior of DOTS clients and client-domain DOTS gateways for each case.

Multi-homed DOTS agents are assumed to make use of the protocols defined in [RFC9132] and [RFC8783]. This document does not require any specific extension to the base DOTS protocols for deploying DOTS in a multi-homed context.

2. Requirements Language

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. Terminology

This document makes use of the terms defined in [RFC8811], [RFC8612], and [RFC4116]. In particular:

Provider-Aggregatable (PA) addresses: globally-unique addresses assigned by a transit provider to a customer. The addresses are considered "aggregatable" because the set of routes corresponding to the PA addresses are usually covered by an aggregate route set corresponding to the address space operated by the transit provider, from which the assignment was made (Section 2 of [RFC4116]).

Provider-Independent (PI) addresses: globally-unique addresses that are not assigned by a transit provider, but are provided by some other organisation, usually a Regional Internet Registry (RIR) (Section 2 of [RFC4116]).

IP indifferently refers to IPv4 or IPv6.

4. Multi-Homing Scenarios

This section describes some multi-homing scenarios that are relevant to DOTS. In the following subsections, only the connections of border routers are shown; internal network topologies are not elaborated.

A multihomed network may enable DOTS for all or a subset of its upstream interconnection links. In such a case, DOTS servers can be explicitly configured or dynamically discovered by a DOTS client using means such as those discussed in [RFC8973]. These DOTS servers can be owned by the upstream provider, managed by a third-party (e.g., mitigation service provider), or a combination thereof.

If a DOTS server is explicitly configured, it is assumed that an interface is also provided to bind the DOTS service to an interconnection link. If no interface is provided, this means that the DOTS server can be reached via any active interface.

This section distinguishes between residential CPEs vs. enterprise CPEs because PI addresses may be used for enterprises while this is not the current practice for residential CPEs.

In the following subsections, all or a subset of interconnection links are associated with DOTS servers.

4.1. Multi-Homed Residential Single CPE

The scenario shown in Figure 2 is characterized as follows:

* The home network is connected to the Internet using one single CPE.

* The CPE is connected to multiple provisioning domains (i.e., both fixed and mobile networks). Provisioning domain (PvD) is explained in [RFC7556].

In a typical deployment scenario, these provisioning domains are owned by the same provider (see Section 1 of [RFC8803]). Such a deployment is meant to seamlessly use both fixed and cellular networks for bonding, faster hand-overs, or better resiliency purposes.

* Each of these provisioning domains assigns IP addresses/prefixes to the CPE and provides additional configuration information such as a list of DNS servers, DNS suffixes associated with the network, default gateway address, and DOTS server's name [RFC8973]. These addresses/prefixes are assumed to be Provider-Aggregatable (PA).

* Because of ingress filtering, packets forwarded by the CPE towards a given provisioning domain must be sent with a source IP address that was assigned by that domain [RFC8043].

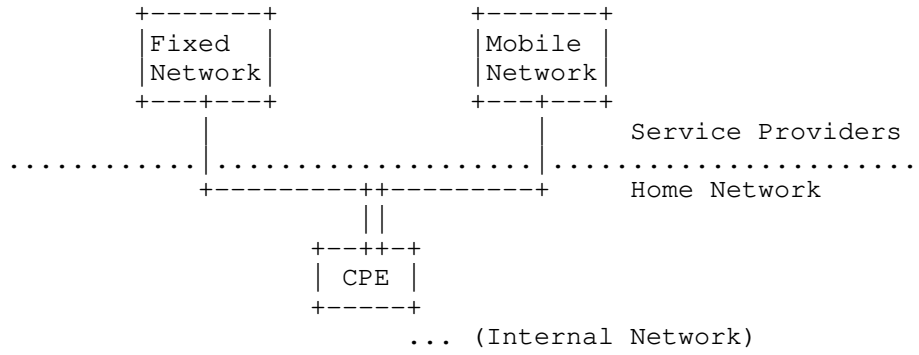

Figure 2: Typical Multi-homed Residential CPE

4.2. Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs

The scenario shown in Figure 3 is characterized as follows:

* The enterprise network is connected to the Internet using a single router.

* That router is connected to multiple provisioning domains managed by distinct administrative entities.

Unlike the previous scenario, two sub-cases can be considered for an enterprise network with regards to assigned addresses:

1. PI addresses/prefixes: The enterprise is the owner of the IP addresses/prefixes; the same address/prefix is then used when establishing communications over any of the provisioning domains.

2. PA addresses/prefixes: Each of the provisioning domains assigns IP addresses/prefixes to the enterprise network. These addresses/prefixes are used when communicating over the provisioning domain that assigned them.

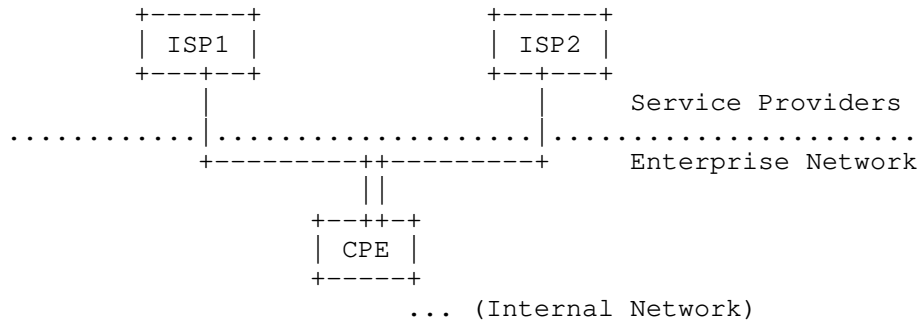

Figure 3: Multi-homed Enterprise Network (Single CPE connected to Multiple Networks)

4.3. Multi-homed Enterprise: Multiple CPEs, Multiple Upstream ISPs

This scenario is similar to the one described in Section 4.2; the main difference is that dedicated routers (CPE1 and CPE2) are used to connect to each provisioning domain.

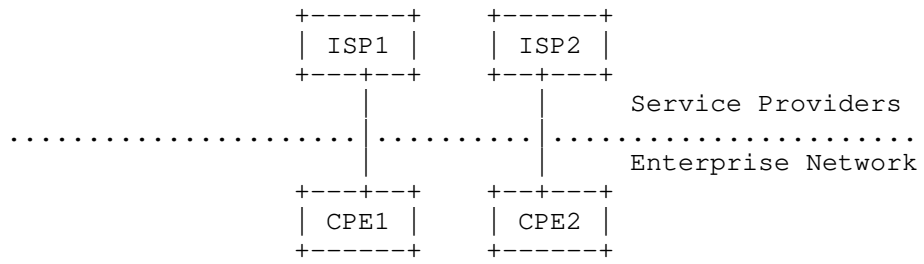

Figure 4: Multi-homed Enterprise Network (Multiple CPEs, Multiple ISPs)

4.4. Multi-homed Enterprise with the Same ISP

This scenario is a variant of Sections 4.2 and 4.3 in which multi-homing is supported by the same ISP (i.e., same provisioning domain).

5. DOTS Multi-homing Deployment Considerations

Table 1 provides some sample, non-exhaustive, deployment schemes to illustrate how DOTS agents may be deployed for each of the scenarios introduced in Section 4.

Scenario	DOTS client	Client-domain DOTS gateway
Residential CPE	CPE	N/A
Single CPE, Multiple provisioning domains	Internal hosts or CPE	CPE
Multiple CPEs, Multiple provisioning domains	Internal hosts or all CPEs (CPE1 and CPE2)	CPEs (CPE1 and CPE2)
Multi-homed enterprise, Single provisioning domain	Internal hosts or all CPEs (CPE1 and CPE2)	CPEs (CPE1 and CPE2)

Table 1: Sample Deployment Cases

These deployment schemes are further discussed in the following subsections.

5.1. Residential CPE

Figure 5 depicts DOTS sessions that need to be established between a DOTS client (C) and two DOTS servers (S1, S2) within the context of the scenario described in Section 4.1. As listed in Table 1, the DOTS client is hosted by the residential CPE.

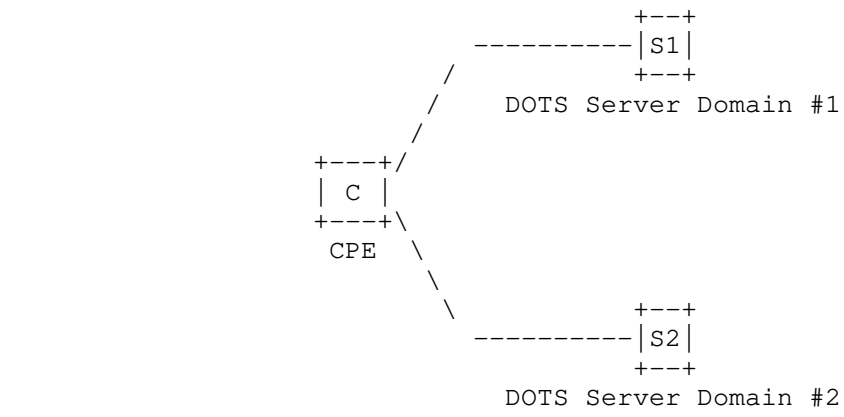

Figure 5: DOTS Associations for a Multihomed Residential CPE

The DOTS client MUST resolve the DOTS server's name provided by each provisioning domain using either the DNS servers learned from the respective provisioning domain or from the DNS servers associated with the interface(s) for which a DOTS server was explicitly configured (Section 4). IPv6-capable DOTS clients MUST use the source address selection algorithm defined in [RFC6724] to select the candidate source addresses to contact each of these DOTS servers. DOTS sessions MUST be established and MUST be maintained with each of the DOTS servers because the mitigation scope of each of these servers is restricted. The DOTS client MUST use the security credentials (a certificate, typically) provided by a provisioning domain to authenticate itself to the DOTS server(s) provided by the same provisioning domain. How such security credentials are provided to the DOTS client is out of the scope of this document. The reader may refer to Section 7.1 of [RFC9132] for more details about DOTS authentication methods.

When conveying a mitigation request to protect the attack target(s), the DOTS client MUST select an available DOTS server whose network has assigned the IP prefixes from which target prefixes/addresses are derived. This implies that if no appropriate DOTS server is found, the DOTS client MUST NOT send the mitigation request to any other available DOTS server.

For example, a mitigation request to protect target resources bound to a PA IP address/prefix cannot be satisfied by a provisioning domain other than the one that owns those addresses/prefixes. Consequently, if a CPE detects a DDoS attack that spreads over all its network attachments, it MUST contact all DOTS servers for mitigation purposes.

The DOTS client MUST be able to associate a DOTS server with each provisioning domain it serves. For example, if the DOTS client is provisioned with S1 using DHCP when attaching to a first network and with S2 using Protocol Configuration Option (PCO) [TS.24008] when attaching to a second network, the DOTS client must record the interface from which a DOTS server was provisioned. A DOTS signaling session to a given DOTS server must be established using the interface from which the DOTS server was provisioned. If a DOTS server is explicitly configured, DOTS signaling with that server must be established via the interfaces that are indicated in the explicit configuration or via any active interface if no interface is configured.

5.2. Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs

Figure 6 illustrates the DOTS sessions that can be established with a client-domain DOTS gateway (hosted within the CPE as per Table 1), which is enabled within the context of the scenario described in Section 4.2. This deployment is characterized as follows:

* One or more DOTS clients are enabled in hosts located in the internal network.

* A client-domain DOTS gateway is enabled to aggregate and then relay the requests towards upstream DOTS servers.

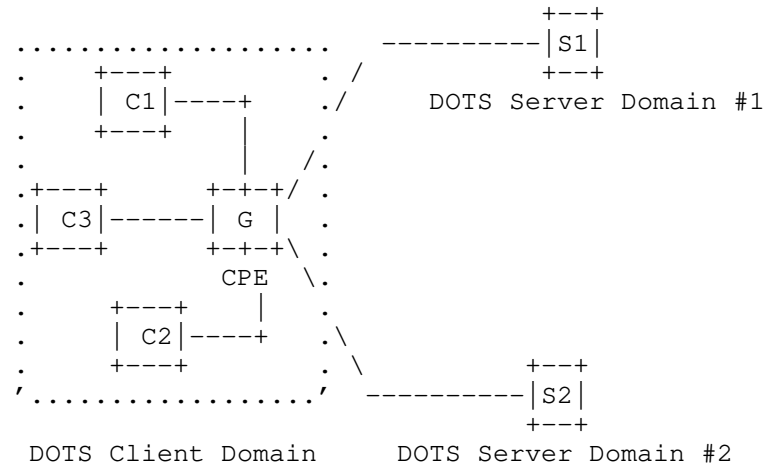

Figure 6: Multiple DOTS Clients, Single DOTS Gateway, Multiple DOTS Servers

When PA addresses/prefixes are in use, the same considerations discussed in Section 5.1 need to be followed by the client-domain DOTS gateway to contact its DOTS server(s). The client-domain DOTS gateways can be reachable from DOTS clients by using a unicast address or an anycast address (Section 3.2.4 of [RFC8811]).

Nevertheless, when PI addresses/prefixes are assigned and absent any policy, the client-domain DOTS gateway SHOULD send mitigation requests to all its DOTS servers. Otherwise, the attack traffic may still be delivered via the ISP that hasn't received the mitigation request.

An alternate deployment model is depicted in Figure 7. This deployment assumes that:

* One or more DOTS clients are enabled in hosts located in the internal network. These DOTS clients may use [RFC8973] to discover their DOTS server(s).

* These DOTS clients communicate directly with upstream DOTS servers.

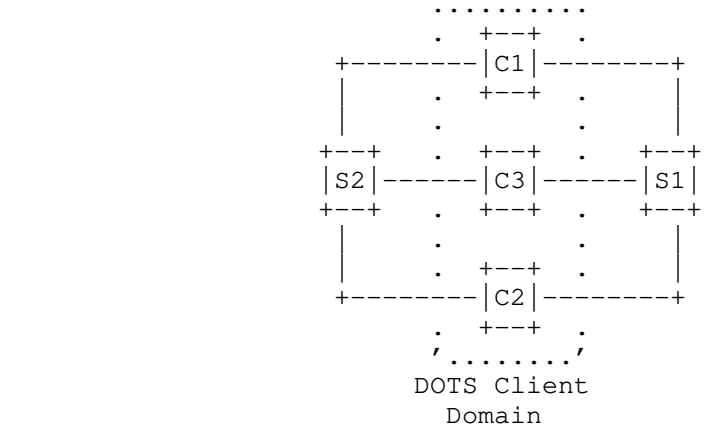

Figure 7: Multiple DOTS Clients, Multiple DOTS Servers

If PI addresses/prefixes are in use, the DOTS client MUST send a mitigation request to all the DOTS servers. The use of the same anycast addresses to reach these DOTS servers is NOT RECOMMENDED. If a well-known anycast address is used to reach multiple DOTS servers, the CPE may not be able to select the appropriate provisioning domain to which the mitigation request should be forwarded. As a consequence, the request may not be forwarded to the appropriate DOTS server.

If PA addresses/prefixes are used, the same considerations discussed in Section 5.1 need to be followed by the DOTS clients. Because DOTS clients are not embedded in the CPE and multiple addresses/prefixes may not be assigned to the DOTS client (typically in an IPv4 context), some issues may arise in how to steer traffic towards the appropriate DOTS server by using the appropriate source IP address. These complications discussed in [RFC4116] are not specific to DOTS.

Another deployment approach is to enable many DOTS clients; each of them is responsible for handling communications with a specific DOTS server (see Figure 8).

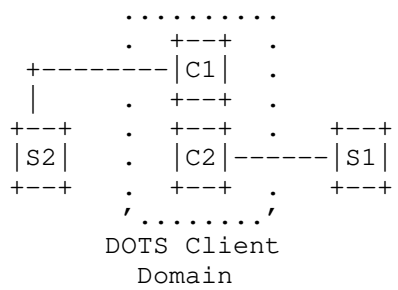

Figure 8: Single Homed DOTS Clients

For both deployments depicted in Figures 7 and 8, each DOTS client SHOULD be provided with policies (e.g., a prefix filter that is used to filter DDoS detection alarms) that will trigger DOTS communications with the DOTS servers. Such policies will help the DOTS client to select the appropriate destination DOTS server. The CPE MUST select the appropriate source IP address when forwarding DOTS messages received from an internal DOTS client.

5.3. Multi-Homed Enterprise: Multiple CPEs, Multiple Upstream ISPs

The deployments depicted in Figures 7 and 8 also apply to the scenario described in Section 4.3. One specific problem for this scenario is to select the appropriate exit router when contacting a given DOTS server.

An alternative deployment scheme is shown in Figure 9:

* DOTS clients are enabled in hosts located in the internal network.

* A client-domain DOTS gateway is enabled in each CPE (CPE1 and CPE2 per Table 1).

* Each of these client-domain DOTS gateways communicates with the DOTS server of the provisioning domain.

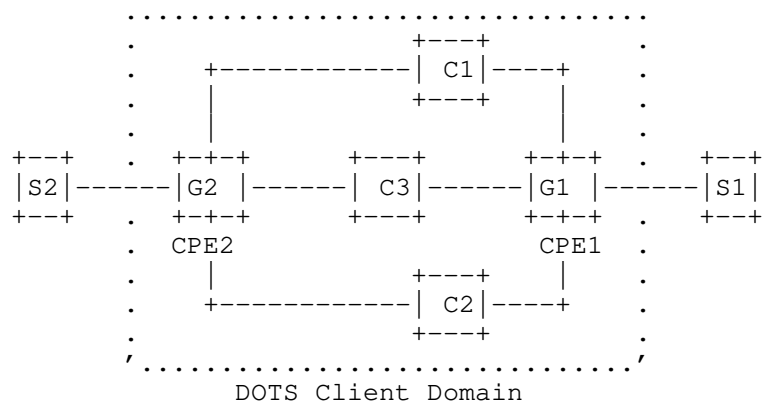

Figure 9: Multiple DOTS Clients, Multiple DOTS Gateways, Multiple DOTS Servers

When PI addresses/prefixes are used, DOTS clients MUST contact all the client-domain DOTS gateways to send a DOTS message. Client-domain DOTS gateways will then relay the request to the DOTS servers as a function of local policy. Note that (same) anycast addresses cannot be used to establish DOTS sessions between DOTS clients and client-domain DOTS gateways because only one DOTS gateway will receive the mitigation request.

When PA addresses/prefixes are used, but no filter rules are provided to DOTS clients, the latter MUST contact all client-domain DOTS gateways simultaneously to send a DOTS message. Upon receipt of a request by a client-domain DOTS gateway, it MUST check whether the request is to be forwarded upstream (if the target IP prefix is managed by the upstream server) or rejected.

When PA addresses/prefixes are used, but specific filter rules are provided to DOTS clients using some means that are out of scope of this document, the clients MUST select the appropriate client-domain DOTS gateway to reach. The use of the same anycast addresses is NOT RECOMMENDED to reach client-domain DOTS gateways.

5.4. Multi-Homed Enterprise: Single ISP

The key difference of the scenario described in Section 4.4 compared to the other scenarios is that multi-homing is provided by the same ISP. Concretely, that ISP can decide to provision the enterprise network with:

* The same DOTS server for all network attachments.

* Distinct DOTS servers for each network attachment. These DOTS servers need to coordinate when a mitigation action is received from the enterprise network.

In both cases, DOTS agents enabled within the enterprise network MAY decide to select one or all network attachments to send DOTS mitigation requests.

6. Security Considerations

A set of security threats related to multihoming are discussed in [RFC4218].

DOTS-related security considerations are discussed in Section 4 of [RFC8811].

DOTS clients should control the information that they share with peer DOTS servers. In particular, if a DOTS client maintains DOTS sessions with specific DOTS servers per interconnection link, the DOTS client SHOULD NOT leak information specific to a given link to DOTS servers on different interconnection links that are not authorized to mitigate attacks for that given link. Whether this constraint is relaxed is deployment-specific and must be subject to explicit consent from the DOTS client domain administrator. How to seek for such consent is implementation- and deployment-specific.

7. IANA Considerations

This document does not require any action from IANA.

8. Acknowledgements

Thanks to Roland Dobbins, Nik Teague, Jon Shallow, Dan Wing, and Christian Jacquenet for sharing their comments on the mailing list.

Thanks to Kirill Kasavchenko for the comments.

Thanks to Kathleen Moriarty for the secdir review, Joel Jaeggli for the opsdir review, Mirja Kuhlewind for the tsvart review, and Dave Thaler for the Intdir review.

Many thanks to Roman Danyliw for the careful AD review.

Thanks to Lars Eggert, Robert Wilton, Paul Wouters, Erik Kline, and Eric Vyncke for the IESG review.

9. References

9.1. Normative References

[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>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <https://www.rfc-editor.org/info/rfc6724>.

[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>.

[RFC8811] Mortensen, A., Ed., Reddy.K, T., Ed., Andreasen, F., Teague, N., and R. Compton, "DDoS Open Threat Signaling (DOTS) Architecture", RFC 8811, DOI 10.17487/RFC8811, August 2020, <https://www.rfc-editor.org/info/rfc8811>.

9.2. Informative References

[RFC3582] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, DOI 10.17487/RFC3582, August 2003, <https://www.rfc-editor.org/info/rfc3582>.

[RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, DOI 10.17487/RFC4116, July 2005, <https://www.rfc-editor.org/info/rfc4116>.

[RFC4218] Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions", RFC 4218, DOI 10.17487/RFC4218, October 2005, <https://www.rfc-editor.org/info/rfc4218>.

[RFC4732] Handley, M., Ed., Rescorla, E., Ed., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, DOI 10.17487/RFC4732, December 2006, <https://www.rfc-editor.org/info/rfc4732>.

[RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <https://www.rfc-editor.org/info/rfc7556>.

[RFC8043] Sarikaya, B. and M. Boucadair, "Source-Address-Dependent Routing and Source Address Selection for IPv6 Hosts: Overview of the Problem Space", RFC 8043, DOI 10.17487/RFC8043, January 2017, <https://www.rfc-editor.org/info/rfc8043>.

[RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <https://www.rfc-editor.org/info/rfc8612>.

[RFC8783] Boucadair, M., Ed. and T. Reddy.K, Ed., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", RFC 8783, DOI 10.17487/RFC8783, May 2020, <https://www.rfc-editor.org/info/rfc8783>.

[RFC8803] Bonaventure, O., Ed., Boucadair, M., Ed., Gundavelli, S., Seo, S., and B. Hesmans, "0-RTT TCP Convert Protocol", RFC 8803, DOI 10.17487/RFC8803, July 2020, <https://www.rfc-editor.org/info/rfc8803>.

[RFC8903] Dobbins, R., Migault, D., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use Cases for DDoS Open Threat Signaling", RFC 8903, DOI 10.17487/RFC8903, May 2021, <https://www.rfc-editor.org/info/rfc8903>.

[RFC8973] Boucadair, M. and T. Reddy.K, "DDoS Open Threat Signaling (DOTS) Agent Discovery", RFC 8973, DOI 10.17487/RFC8973, January 2021, <https://www.rfc-editor.org/info/rfc8973>.

[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>.

[TS.24008] 3GPP, "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3 (Release 16)", December 2019, <http://www.3gpp.org/DynaReport/24008.htm>.

Authors' Addresses

Mohamed Boucadair
Orange
35000 Rennes
France
Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy.K
Akamai
Embassy Golf Link Business Park
Bangalore 560071
Karnataka
India
Email: kondtir@gmail.com

Wei Pan
Huawei Technologies
Email: william.panwei@huawei.com