

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2022

M. Boucadair
Orange
T. Reddy
McAfee
W. Pan
Huawei Technologies
July 6, 2021

Multi-homing Deployment Considerations for Distributed-Denial-of-Service
Open Threat Signaling (DOTS)
[draft-ietf-dots-multihoming-07](#)

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/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 January 7, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4](#).e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Requirements Language	4
3.	Terminology	4
4.	Multi-Homing Scenarios	5
4.1.	Multi-Homed Residential Single CPE	5
4.2.	Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs	6
4.3.	Multi-homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	7
4.4.	Multi-homed Enterprise with the Same ISP	7
5.	DOTS Multi-homing Deployment Considerations	8
5.1.	Residential CPE	8
5.2.	Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs	9
5.3.	Multi-Homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	12
5.4.	Multi-Homed Enterprise: Single ISP	13
6.	Security Considerations	13
7.	IANA Considerations	14
8.	Acknowledgements	14
9.	References	14
9.1.	Normative References	14
9.2.	Informative References	14
	Authors' Addresses	15

[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. It 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.
 - * Blindly forking 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:

- o Identify and extract viable deployment candidates from [\[RFC8903\]](#).
- o 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
- o Describe the recommended behavior of DOTS clients and gateways for each case.

Multi-homed DOTS agents are assumed to make use of the protocols defined in [\[I-D.ietf-dots-rfc8782-bis\]](#) and [\[RFC8783\]](#); no specific extension is required 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 [\[RFC8811\]](#) and [\[RFC4116\]](#). In particular:

Provider-Aggregatable (PA) addresses are 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 are globally-unique addresses which 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.

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.

4.1. Multi-Homed Residential Single CPE

The scenario shown in Figure 2 is characterized as follows:

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

- o 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).
- o 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](#)].

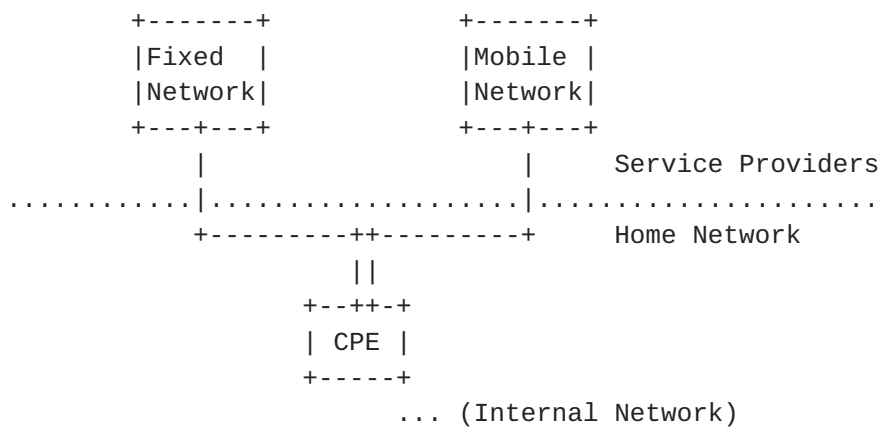


Figure 2: Typical Multi-homed Residential CPE

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

The scenario shown in Figure 3 is characterized as follows:

- o The enterprise network is connected to the Internet using a single router.
- o That router is connected to multiple provisioning domains (i.e., 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.




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 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 [Section 4.2](#) and [Section 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	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 (rtr1 and rtr2)	CPEs (rtr1 and rtr2)
Multi-homed enterprise, Single provisioning domain	Internal hosts or all CPEs (rtr1 and rtr2)	CPEs (rtr1 and rtr2)

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](#).

For each provisioning domain, the DOTS client MUST resolve the DOTS server's name provided by a provisioning domain ([\[RFC8973\]](#)) using the DNS servers learned from the respective provisioning domain. 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 SHOULD use the certificate provisioned by a provisioning domain to authenticate itself to the DOTS server(s) provided by the same provisioning domain.

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 both DOTS servers for mitigation purposes.

The DOTS client MUST be able to associate a DOTS server with each provisioning domain. 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) when attaching to a second network, the DOTS client must record the interface from which a DOTS server was provisioned. DOTS signaling session to a given DOTS server must be established using the interface from which the DOTS server was provisioned.




Figure 5: DOTS Associations for a Multihomed Residential CPE

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

Figure 6 illustrates a first set of DOTS associations that can be established with a DOTS gateway, which is enabled within the context of the scenario described in [Section 4.2](#). This deployment is characterized as follows:

- o One of more DOTS clients are enabled in hosts located in the internal network.

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

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

Nevertheless, when PI addresses/prefixes are assigned, the DOTS gateway MUST send mitigation requests to all its DOTS servers. Otherwise, the attack traffic may still be delivered via the ISP which hasn't received the mitigation request.




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

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

- o 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).
- o These DOTS clients communicate directly with upstream 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 anycast addresses to reach the DOTS servers is NOT RECOMMENDED.

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.




Figure 7: Multiple DOTS Clients, Multiple DOTS Servers

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

Each DOTS client SHOULD be provided with policies (e.g., a prefix filter that will be against 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. If anycast addresses are used to reach 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.

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:

- o DOTS clients are enabled in hosts located in the internal network.
- o A DOTS gateway is enabled in each CPE (rtr1, rtr2).
- o Each of these DOTS gateways communicates with the DOTS server of the provisioning domain.

When PI addresses/prefixes are used, DOTS clients **MUST** contact all the DOTS gateways to send a DOTS message. DOTS gateways will then relay the request to the DOTS server. Note that the use of anycast addresses is **NOT RECOMMENDED** to establish DOTS sessions between DOTS clients and DOTS gateways.

When PA addresses/prefixes are used, but no filter rules are provided to DOTS clients, the latter **MUST** contact all DOTS gateways simultaneously to send a DOTS message. Upon receipt of a request by a 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 DOTS gateway to reach. The use of anycast addresses is **NOT RECOMMENDED** to reach DOTS gateways.




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

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:

- o The same DOTS server for all network attachments.
- o 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

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 associations 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.

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

- [I-D.ietf-dots-rfc8782-bis] Boucadair, M., Shallow, J., and T. Reddy, K., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", [draft-ietf-dots-rfc8782-bis-06](#) (work in progress), March 2021.
- [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>>.
- [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>>.
- [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>>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: TirumaleswarReddy_Konda@McAfee.com

Wei Pan
Huawei Technologies

Email: william.panwei@huawei.com