

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 10, 2019

M. Boucadair
Orange
T. Reddy
McAfee
P. Patil
Cisco
October 7, 2018

Distributed-Denial-of-Service Open Threat Signaling (DOTS) Server
Discovery
draft-boucadair-dots-server-discovery-05

Abstract

It may not be possible for a network to determine the cause for an attack, but instead just realize that some resources seem to be under attack. To fill that gap, Distributed-Denial-of-Service Open Threat Signaling (DOTS) allows a network to inform a DOTS server that it is under a potential attack so that appropriate mitigation actions are undertaken.

This document specifies mechanisms to configure nodes with DOTS servers.

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 April 10, 2019.

Copyright Notice

Copyright (c) 2018 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	3
2. Requirements Language	4
3. Terminology	4
4. Why Multiple Discovery Mechanisms?	5
5. Discovery Procedure	7
6. Resolution	8
7. Discovery using Service Resolution	10
7.1. Retrieving Domain Name	10
7.1.1. DHCP	10
8. DNS Service Discovery	11
8.1. DNS-SD	11
8.2. mDNS	11
9. DHCP Options for DOTS	11
9.1. DHCPv6 DOTS Options	12
9.1.1. Format of DOTS Reference Identifier Option	12
9.1.2. Format Format of DOTS Address Option	13
9.1.3. DHCPv6 Client Behavior	13
9.2. DHCPv4 DOTS Options	14
9.2.1. Format of DOTS Reference Identifier Option	14
9.2.2. Format Format of DOTS Address Option	15
9.2.3. DHCPv4 Client Behavior	16
10. Anycast	17
11. Security Considerations	17
11.1. DHCP	18
11.2. Service Resolution	18
11.3. DNS Service Discovery	18
11.4. Anycast	18
12. IANA Considerations	19
12.1. DHCPv6 Option	19
12.2. DHCPv4 Option	19
12.3. Application Service & Application Protocol Tags	19
12.3.1. DOTS Application Service Tag Registration	19
12.3.2. signal.udp Application Protocol Tag Registration	20
12.3.3. signal.tcp Application Protocol Tag Registration	20
12.3.4. data.tcp Application Protocol Tag Registration	20
12.4. IPv4 Anycast	20

12.5. IPv6 Anycast	21
13. Acknowledgements	21
14. References	22
14.1. Normative References	22
14.2. Informative References	23
Authors' Addresses	24

1. Introduction

In many deployments, it may not be possible for a network to determine the cause for a distributed Denial-of-Service (DoS) attack [RFC4732], but instead just realize that some resources seem to be under attack. To fill that gap, the IETF is specifying an architecture, called DDoS Open Threat Signaling (DOTS) [I-D.ietf-dots-architecture], in which a DOTS client can inform a DOTS server that the network is under a potential attack and that appropriate mitigation actions are required. Indeed, because the lack of a common method to coordinate a real-time response among involved actors and network domains inhibits the effectiveness of DDoS attack mitigation, DOTS protocol is meant to carry requests for DDoS attack mitigation, thereby reducing the impact of an attack and leading to more efficient defensive actions. [I-D.ietf-dots-use-cases] identifies a set of scenarios for DOTS.

The basic high-level DOTS architecture is illustrated in Figure 1 ([I-D.ietf-dots-architecture]):




Figure 1: Basic DOTS Architecture

[I-D.ietf-dots-architecture] specifies that the DOTS client may be provided with a list of DOTS servers; each 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 addresses. The logic for connecting to one or multiple IP addresses is out of scope of this document.

This document specifies methods for DOTS clients to discover their DOTS server(s). The rationale for specifying multiple discovery mechanisms is discussed in Section 4.

Considerations for the selection of DOTS server(s) by multi-homed DOTS clients is out of scope; the reader should refer to [I-D.boucadair-dots-multihoming] for more details.

Likewise, happy eyeballs considerations for DOTS are out of scope. The reader should refer to Section 4 of [I-D.ietf-dots-signal-channel].

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 following terms:

- o DDoS: A distributed Denial-of-Service attack, in which traffic originating from multiple sources are directed at a target on a network. DDoS attacks are intended to cause a negative impact on the availability of servers, services, applications, and/or other functionality of an attack target.
- o DHCP refers to both DHCPv4 [RFC2131] and DHCPv6 [RFC3315].
- o DHCP client denotes a node that initiates requests to obtain configuration parameters from one or more DHCP servers.
- o DHCP server refers to a node that responds to requests from DHCP clients.
- o DOTS client: A DOTS-aware software module responsible for requesting attack response coordination with other DOTS-aware elements.
- o DOTS server: A DOTS-aware software module handling and responding to messages from DOTS clients. The DOTS server should enable mitigation on behalf of the DOTS client, if requested, by communicating the DOTS client's request to the mitigator and returning selected mitigator feedback to the requesting DOTS client. A DOTS server may also be a mitigator.
- o DOTS gateway: A DOTS-aware software module that is logically equivalent to a DOTS client back-to-back with a DOTS server.

Furthermore, the reader should be familiar with other terms defined in [I-D.ietf-dots-architecture] and [RFC3958].

4. Why Multiple Discovery Mechanisms?

It is tempting to specify one single discovery mechanism for DOTS. Nevertheless, the analysis of the various use cases sketched in [I-D.ietf-dots-use-cases] reveals that it is unlikely that one single discovery method can be suitable for all the sample deployments (Table 1). Concretely:

- o Some of the use cases may allow DOTS clients to have direct communications with upstream DOTS servers; that is no DOTS gateway is involved. Leveraging on existing features that do not require specific feature on the node embedding the DOTS client may ease DOTS deployment. Typically, the use of Straightforward-Naming Authority Pointer (S-NAPTR) lookups [RFC3958] allows the DOTS server administrators provision the preferred DOTS signal channel transport protocol between the DOTS client and the DOTS server and allows the DOTS client to discover this preference.
- o Resolving a DOTS server domain name offered by the upstream transit provider provisioned to a DOTS client into IP address(es) require the use of the appropriate DNS resolvers; otherwise, resolving those names will fail. The use of protocols such as DHCP does allow to associate provisioned DOTS server domain names with a list of DNS servers to be used for name resolution.
- o The upstream network provider is not the DDoS mitigation provider for some of these use cases. The use of anycast is not appropriate for this use case, in particular. It is safe to assume that for such deployments, the DOTS server(s) domain name is provided during the service subscription (i.e., manual/local configuration).
- o Multiple DOTS clients may be enabled within a network (e.g., enterprise network). Automatic means to discover DOTS servers in a deterministic manner are interesting from an operational standpoint.
- o Some of the use cases may involve a DOTS gateway that is responsible for forking requests received from DOTS clients to upstream DOTS servers or for selecting the appropriate DOTS server. Particularly, the use of anycast may simplify the operations within the enterprise network to discover a DOTS gateway, if the enterprise network is single-homed.
- o Many use cases discussed in [I-D.ietf-dots-use-cases] do involve a CPE device. Multiple CPEs, connected to distinct network providers may even be considered. It is intuitive to leverage on existing mechanisms such as discovery using service resolution or

DHCP or anycast to provision the CPE acting as a DOTS client with the DOTS server(s).

Use Case	Requires a CPE	The Network Provider is also the DDoS Mitigation Provider
End-customer with single or multiple upstream transit provider(s) offering DDoS mitigation services	Yes (Intelligent DDoS mitigation system (IDMS) acting as a DOTS client may be co-located on the CPE)	Yes
End-customer with an overlay DDoS mitigation managed security service provider (MSSP)	Yes (DDoS Detector acting as a DOTS client may be co-located on the CPE)	No
End-customer operating an application or service with an integrated DOTS client	Yes (CPE may act as a DOTS gateway)	Yes/No
End-customer operating a CPE network infrastructure device with an integrated DOTS client	Yes (CPE acts as a DOTS client)	Yes
Suppression of outbound DDoS traffic originating from a consumer broadband access network	Yes (CPE acts as a DOTS server)	Yes
DDoS Orchestration	No	N/A

Table 1: Summary of DOTS Use Cases

Consequently, this document describes the following mechanisms for discovery:

- o A resolution mechanism based on straightforward Naming Authority Pointer (S-NAPTR) resource records in the Domain Name System (DNS).
- o DNS Service Discovery.
- o Discovery using DHCP Options.
- o A mechanism based on anycast address for DOTS usage.

5. Discovery Procedure

A key point in the deployment of DOTS is the ability of network operators to be able to configure DOTS clients with the correct server information consistently. To accomplish this, operators will need a consistent set of ways in which DOTS clients can discover this information, and a consistent priority among these options. If some devices prefer manual configuration over DNS discovery, while others prefer DNS discovery over manual configuration, the result will be a process of "whack-a-mole", where the operator must find devices that are using the wrong DOTS server, determine how to ensure the devices are configured properly, and then reconfigure the device through the preferred method.

All DOTS clients MUST support at least one of the four mechanisms below to determine a DOTS server list. All DOTS clients SHOULD implement all four, or as many as are practical for any specific device, of these ways to discover DOTS servers, in order to facilitate the deployment of DOTS in large scale environments:

1. Explicit configuration:

- * Local/Manual configuration: A DOTS client, will learn the DOTS server(s) by means of local or manual DOTS configuration (i.e., DOTS servers configured at the system level). Configuration discovered from a DOTS client application is considered as local configuration. An implementation may give the user an opportunity (e.g., by means of configuration file options or menu items) to specify DOTS server(s) for each address family. These MAY be specified either as IP addresses or the DNS name of a DOTS server. When only DOTS server' IP addresses are configured, a reference identifier must also be configured for authentication purposes.
- * Automatic configuration (e.g., DHCP, an automation system): The DOTS client attempts to discover DOTS server(s) names and/or addresses from DHCP, as described in Section 9.

2. Service Resolution : The DOTS client attempts to discover DOTS server name(s) using service resolution, as specified in Section 7.
3. DNS SD: DNS Service Discovery. The DOTS client attempts to discover DOTS server name(s) using DNS service discovery, as specified in Section 8.
4. Anycast : Send DOTS request to establish a DOTS session with the assigned DOTS server anycast address for each combination of interface and address family.

Some of these mechanisms imply the use of DNS to resolve the IP address of the DOTS server, while others imply the IP address of the relevant DOTS server is obtained directly. Implementation options may vary on a per device basis, as some devices may not have DNS capabilities and/or proper configuration.

Clients will prefer information received from the discovery methods in the order listed.

On hosts with more than one interface or address family (IPv4/v6), the DOTS server discovery procedure has to be performed for each combination of interface and address family. A client MAY choose to perform the discovery procedure only for a desired interface/address combination if the client does not wish to discover a DOTS server for all combinations of interface and address family.

The above procedure MUST also be followed by a DOTS gateway.

6. Resolution

Once the DOTS client has retrieved client's DNS domain or discovered the DOTS server name that needs to be resolved, an S-NAPTR lookup with 'DOTS' application service and the desired protocol tag is made to obtain information necessary to connect to the authoritative DOTS server within the given domain.

This specification defines "DOTS" as an application service tag (Section 12.3.1) and "signal.udp" (Section 12.3.2), "signal.tcp" (Section 12.3.3), and "data.tcp" (Section 12.3.4) as application protocol tags.

In the example below, for domain 'example.net', the resolution algorithm will result in IP address(es), port, tag and protocol tuples as follows:

```
example.net.
IN NAPTR 100 10 "" DOTS:signal.udp "" signal.example.net.
IN NAPTR 200 10 "" DOTS:signal.tcp "" signal.example.net.
IN NAPTR 300 10 "" DOTS:data.tcp "" data.example.net.

signal.example.net.
IN NAPTR 100 10 S DOTS:signal.udp "" _dots._signal._udp.example.net.
IN NAPTR 200 10 S DOTS:signal.tcp "" _dots._signal._tcp.example.net.

data.example.net.
IN NAPTR 100 10 S DOTS:data.tcp "" _dots._data._tcp.example.net.

_dots._signal._udp.example.net.
IN SRV 0 0 5000 a.example.net.

_dots._signal._tcp.example.net.
IN SRV 0 0 5001 a.example.net.

_dots._data._tcp.example.net.
IN SRV 0 0 5002 a.example.net.

a.example.net.
IN AAAA 2001:db8::1
```

Order	Protocol	IP address	Port	Tag
1	UDP	2001:db8::1	5000	Signal
2	TCP	2001:db8::1	5001	Signal
3	TCP	2001:db8::1	5002	Data

If no DOTS-specific S-NAPTR records can be retrieved, the discovery procedure fails for this domain name (and the corresponding interface and IP protocol version). If more domain names are known, the discovery procedure MAY perform the corresponding S-NAPTR lookups immediately. However, before retrying a lookup that has failed, a DOTS client MUST wait a time period that is appropriate for the encountered error (e.g., NXDOMAIN, timeout, etc.).

7. Discovery using Service Resolution

This mechanism is performed in two steps:

1. A DNS domain name is retrieved for each combination of interface and address family.
2. Retrieved DNS domain names are then used for S-NAPTR lookups. Further DNS lookups may be necessary to determine DOTS server IP address(es).

7.1. Retrieving Domain Name

A DOTS client has to determine the domain in which it is located. The following section describes the means to obtain the domain name from DHCP. Other means of retrieving domain names may be used, which are outside the scope of this document, e.g., local configuration.

Implementations MAY allow the user to specify a default name that is used, if no specific name has been configured.

7.1.1. DHCP

DHCP can be used to determine the domain name related to an interface's point of network attachment. Network operators may provide the domain name to be used for service discovery within an access network using DHCP. Sections 3.2 and 3.3 of [RFC5986] define DHCP IPv4 and IPv6 access network domain name options, `OPTION_V4_ACCESS_DOMAIN` and `OPTION_V6_ACCESS_DOMAIN` respectively, to identify a domain name that is suitable for service discovery within the access network.

For IPv4, the discovery procedure MUST request the access network domain name option in a Parameter Request List option, as described in [RFC2131]. [RFC2132] defines the DHCP IPv4 domain name option; while this option is less suitable, a client MAY request for it if the access network domain name defined in [RFC5986] is not available.

For IPv6, the discovery procedure MUST request for the access network domain name option in an Options Request Option (ORO) within an Information-request message, as described in [RFC3315].

If neither option can be retrieved the procedure fails for this interface. If a result can be retrieved it will be used as an input for S-NAPTR resolution discussed in Section 6.

8. DNS Service Discovery

DNS-based Service Discovery (DNS-SD) [RFC6763] and Multicast DNS (mDNS) [RFC6762] provide generic solutions for discovering services. DNS-SD/mDNS define a set of naming rules for certain DNS record types that they use for advertising and discovering services.

8.1. DNS-SD

Section 4.1 of [RFC6763] specifies that a service instance name in DNS-SD has the following structure:

```
<Instance> . <Service> . <Domain>
```

The <Domain> portion specifies the DNS sub-domain where the service instance is registered. It may be "local.", indicating the mDNS local domain, or it may be a conventional domain name such as "example.com."

The <Service> portion of the DOTS service instance name MUST be "_dots._signal._udp" or "_dots._signal._tcp" or "_dots._data._tcp".

8.2. mDNS

A DOTS client can proactively discover DOTS servers being advertised in the site by multicasting a PTR query to one or all of the following:

- o "_dots._signal._udp.local."
- o "_dots._signal._tcp.local."
- o "_dots._data._tcp.local."

A DOTS server can send out gratuitous multicast DNS answer packets whenever it starts up, wakes from sleep, or detects a change in network configuration. DOTS clients receive these gratuitous packets and cache information contained in it.

9. DHCP Options for DOTS

As reported in Section 1.7.2 of [RFC6125]:

"few certification authorities issue server certificates based on IP addresses, but preliminary evidence indicates that such certificates are a very small percentage (less than 1%) of issued certificates".

In order to allow for PKIX-based authentication between a DOTS client and server while accommodating for the current best practices for issuing certificates, this document allows for configuring names to DOTS clients. These names can be used for two purposes: to retrieve the list of IP addresses of a DOTS server or to be presented as a reference identifier for authentication purposes.

Defining the option to include a list of IP addresses would avoid a dependency on an underlying name resolution, but that design requires to also supply a name for PKIX-based authentication purposes.

9.1. DHCPv6 DOTS Options

9.1.1. Format of DOTS Reference Identifier Option

The DHCPv6 DOTS option is used to configure a name of the DOTS server. The format of this option is shown in Figure 2.




Figure 2: DHCPv6 DOTS Reference Identifier option

The fields of the option shown in Figure 2 are as follows:

- o Option-code: OPTION_V6_DOTS_RI (TBA1, see Section 12.1)
- o Option-length: Length of the dots-server-name field in octets.
- o dots-server-name: A fully qualified domain name of the DOTS server. This field is formatted as specified in Section 8 of [RFC3315].

An example of the dots-server-name encoding is shown in Figure 3. This example conveys the FQDN "dots.example.com."




Figure 3: An example of the dots-server-name encoding

9.1.2. Format Format of DOTS Address Option

The DHCPv6 DOTS option can be used to configure a list of IPv6 addresses of a DOTS server. The format of this option is shown in Figure 4.




Figure 4: DHCPv6 DOTS Address option

The fields of the option shown in Figure 4 are as follows:

- o Option-code: OPTION_V6_DOTS_ADDRESS (TBA2, see Section 12.1)
- o Option-length: Length of the 'DOTS ipv6-address(es)' field in octets. MUST be a multiple of 16.
- o DOTS ipv6-address: Includes one or more IPv6 addresses [RFC4291] of the DOTS server to be used by the DOTS client.

Note, IPv4-mapped IPv6 addresses (Section 2.5.5.2 of [RFC4291]) are allowed to be included in this option.

To return more than one DOTS servers to the requesting DHCPv6 client, the DHCPv6 server returns multiple instances of OPTION_V6_DOTS.

9.1.3. DHCPv6 Client Behavior

DHCP clients MAY request options OPTION_V6_DOTS_RI and OPTION_V6_DOTS_ADDRESS, as defined in [RFC3315], Sections 17.1.1, 18.1.1, 18.1.3, 18.1.4, 18.1.5, and 22.7. As a convenience to the reader, it is mentioned here that the DHCP client includes the requested option codes in the Option Request Option.

If the DHCP client receives more than one instance of `OPTION_V6_DOTS_RI` (resp. `OPTION_V6_DOTS_ADDRESS`) option, it MUST use only the first instance of that option.

If the DHCP client receives both `OPTION_V6_DOTS_RI` and `OPTION_V6_DOTS_ADDRESS`, the content of `OPTION_V6_DOTS_RI` is used as reference identifier for authentication purposes (e.g., PKIX [RFC6125]), while the addresses included in `OPTION_V6_DOTS_ADDRESS` are used to reach the DOTS server. In other words, the name conveyed in `OPTION_V6_DOTS_RI` MUST NOT be passed to underlying resolution library in the presence of `OPTION_V6_DOTS_ADDRESS` in a response.


If the DHCP client receives `OPTION_V6_DOTS_RI` only, but `OPTION_V6_DOTS_RI` option contains more than one name, as distinguished by the presence of multiple root labels, the DHCP client MUST use only the first name. Once the name is validated (Section 8 of [RFC3315]), the name is passed to a name resolution library. Moreover, that name is also used as a reference identifier for authentication purposes.

If the DHCP client receives `OPTION_V6_DOTS_ADDRESS` only, the address(es) included in `OPTION_V6_DOTS_ADDRESS` is used to reach the DOTS server. In addition, these addresses can be used as identifiers for authentication.

9.2. DHCPv4 DOTS Options

9.2.1. Format of DOTS Reference Identifier Option

The DHCPv4 DOTS option is used to configure a name of the DOTS server. The format of this option is illustrated in Figure 5.



The values `s1`, `s2`, `s3`, etc. represent the domain name labels in the domain name encoding.

Figure 5: DHCPv4 DOTS Reference Identifier option

The fields of the option shown in Figure 5 are as follows:

- o Code: `OPTION_V4_DOTS_RI` (TBA3, see Section 12.2);
- o Length: Includes the length of the "DOTS server name" field in octets; the maximum length is 255 octets.

- o DOTS server name: The domain name of the DOTS server. This field is formatted as specified in Section 8 of [RFC3315].

9.2.2. Format of DOTS Address Option

The DHCPv4 DOTS option can be used to configure a list of IPv4 addresses of a DOTS server. The format of this option is illustrated in Figure 6.





Figure 6: DHCPv4 DOTS Address option

The fields of the option shown in Figure 6 are as follows:

- o Code: OPTION_V4_DOTS_ADDRESS (TBA4, see Section 12.2);
- o Length: Length of all included data in octets. The minimum length is 5.
- o List-Length: Length of the "List of DOTS IPv4 Addresses" field in octets; MUST be a multiple of 4.
- o List of DOTS IPv4 Addresses: Contains one or more IPv4 addresses of the DOTS server to be used by the DOTS client. The format of this field is shown in Figure 7.
- o OPTION_V4_DOTS can include multiple lists of DOTS IPv4 addresses; each list is treated separately as it corresponds to a given DOTS server.

When several lists of DOTS IPv4 addresses are to be included, "List-Length" and "DOTS IPv4 Addresses" fields are repeated.



This format assumes that an IPv4 address is encoded as a1.a2.a3.a4.

Figure 7: Format of the List of DOTS IPv4 Addresses

OPTION_V4_DOTS is a concatenation-requiring option. As such, the mechanism specified in [RFC3396] MUST be used if OPTION_V4_DOTS exceeds the maximum DHCPv4 option size of 255 octets.

9.2.3. DHCPv4 Client Behavior

To discover a DOTS server, the DHCPv4 client MUST include both OPTION_V4_DOTS_RI and OPTION_V4_DOTS_ADDRESS in a Parameter Request List Option [RFC2132].

If the DHCP client receives more than one instance of OPTION_V4_DOTS_RI (resp. OPTION_V4_DOTS_ADDRESS) option, it MUST use only the first instance of that option.

If the DHCP client receives both OPTION_V4_DOTS_RI and OPTION_V4_DOTS_ADDRESS, the content of OPTION_V4_DOTS_RI is used as reference identifier for authentication purposes, while the addresses included in OPTION_V4_DOTS_ADDRESS are used to reach the DOTS server. In other words, the name conveyed in OPTION_V4_DOTS_RI MUST NOT be passed to underlying resolution library in the presence of OPTION_V4_DOTS_ADDRESS in a response.

If the DHCP client receives OPTION_V4_DOTS_RI only, but OPTION_V4_DOTS_RI option contains more than one name, as distinguished by the presence of multiple root labels, the DHCP client MUST use only the first name. Once the name is validated (Section 8 of [RFC3315]), the name is passed to a name resolution library. Moreover, that name is also used as a reference identifier for authentication purposes.

If the DHCP client receives OPTION_V4_DOTS_ADDRESS only, the address(es) included in OPTION_V4_DOTS_ADDRESS is used to reach the DOTS server. In addition, these addresses can be used as identifiers for authentication.

10. Anycast

IP anycast can also be used for DOTS service discovery. A packet sent to an anycast address is delivered to the 'topologically nearest' network interface with the anycast address.

When a DOTS client requires DOTS services, it attempts to establish a signaling session with the assigned anycast address(es) defined in Sections 12.4 and 12.5. A DOTS server, that receives a DOTS request with an anycast address, SHOULD redirect the DOTS client to the appropriate DOTS unicast server(s) using the mechanism described in Section 5.5 of [I-D.ietf-dots-signal-channel], unless it is configured otherwise. Indeed, a DOTS server SHOULD be configurable to maintain all DOTS communications using anycast. DOTS redirect is not made mandatory because the use of anycast is not problematic for some deployment scenarios such as an enterprise network deploying one single DOTS gateway connected to one single network provider.

[I-D.boucadair-dots-multihoming] identifies a set of deployment schemes in which the use of anycast is not recommended.

11. Security Considerations

DOTS-related security considerations are discussed in Section 4 of [I-D.ietf-dots-architecture] is to be considered. DOTS agents must authenticate each other using (D)TLS before a DOTS session is considered valid.

If the DOTS client is explicitly configured with DOTS server(s) then the DOTS client can also be explicitly configured with credentials to authenticate the DOTS server.

The CPE device acting as a DOTS client MAY use Bootstrapping Remote Secure Key Infrastructures (BRSKI) discussed in [I-D.ietf-anima-bootstrapping-keyinfra] to automatically bootstrap using the vendor installed X.509 certificate, in combination with a domain registrar provided by the upstream transit provider and vendor's authorizing service. The CPE device authenticates to the upstream transit provider using the vendor installed X.509 certificate and the upstream transit provider validates the vendor installed certificate on the CPE device using the Manufacturer Authorized Signing Authority (MASA) service. If authentication is successful then the CPE device can request and get a voucher from the MASA service via the domain registrar. The voucher is signed by the MASA service and includes the upstream transit provider's trust anchor certificate. The CPE device validates the signed voucher using the manufacturer installed trust anchor associated with the vendor's selected MASA service and stores the upstream transit

provider's trust anchor certificate. The CPE device then uses Enrollment over Secure Transport (EST) [RFC7030] for certificate enrollment (Section 3.8 in [I-D.ietf-anima-bootstrapping-keyinfra]). The DOTS client on the CPE device can authenticate to the DOTS server using the certificate provisioned by the EST server and the DOTS client can validate the DOTS server certificate using the upstream transit provider's trust anchor certificate it had received in the voucher.

11.1. DHCP

The security considerations in [RFC2131] and [RFC3315] are to be considered.

11.2. Service Resolution

The primary attack against the methods described in Section 7 is one that would lead to impersonation of a DOTS server. An attacker could attempt to compromise the S-NAPTR resolution. The use of mutual authentication makes it difficult to redirect a DOTS client to an illegitimate DOTS server.

11.3. DNS Service Discovery

Since DNS-SD is just a specification for how to name and use records in the existing DNS system, it has no specific additional security requirements over and above those that already apply to DNS queries and DNS updates. For DNS queries, DNS Security Extensions (DNSSEC) [RFC4033] SHOULD be used where the authenticity of information is important. For DNS updates, secure updates [RFC2136][RFC3007] SHOULD generally be used to control which clients have permission to update DNS records.

For mDNS, in addition to what has been described above, a principal security threat is a security threat inherent to IP multicast routing and any application that runs on it. A rogue system can advertise that it is a DOTS server. Discovery of such rogue systems as DOTS servers, in itself, is not a security threat if the DOTS client authenticates the discovered DOTS servers.

11.4. Anycast

Anycast-related security considerations are discussed in [RFC4786] and [RFC7094].

12. IANA Considerations

IANA is requested to allocate the SRV service name of "_dots._signal" for DOTS signal channel over UDP or TCP, and the service name of "_dots._data" for DOTS data channel over TCP.

12.1. DHCPv6 Option

IANA is requested to assign the following new DHCPv6 Option Code in the registry maintained in <http://www.iana.org/assignments/dhcpv6-parameters/>:

Option Name	Value
OPTION_V6_DOTS_RI	TBA1
OPTION_V6_DOTS_ADDRESS	TBA2

12.2. DHCPv4 Option

IANA is requested to assign the following new DHCPv4 Option Code in the registry maintained in <http://www.iana.org/assignments/bootp-dhcp-parameters/>:

Option Name	Value	Data length	Meaning
OPTION_V4_DOTS_RI	TBA3	Variable; the maximum length is 255 octets.	Includes the name of the DOTS server.
OPTION_V4_DOTS_ADDRESS	TBA4	Variable; the minimum length is 5.	Includes one or multiple lists of DOTS IP addresses; each list is treated as a separate DOTS server.

12.3. Application Service & Application Protocol Tags

This document requests IANA to make the following allocations from the registry available at: <https://www.iana.org/assignments/s-naptr-parameters/s-naptr-parameters.xhtml>.

12.3.1. DOTS Application Service Tag Registration

- o Application Protocol Tag: DOTS
- o Intended Usage: See Section 6
- o Security Considerations: See Section 11

- o Contact Information: <one of the authors>

12.3.2. signal.udp Application Protocol Tag Registration

- o Application Protocol Tag: signal.udp
- o Intended Usage: See Section 6
- o Security Considerations: See Section 11
- o Contact Information: <one of the authors>

12.3.3. signal.tcp Application Protocol Tag Registration

- o Application Protocol Tag: signal.tcp
- o Intended Usage: See Section 6
- o Security Considerations: See Section 11
- o Contact Information: <one of the authors>

12.3.4. data.tcp Application Protocol Tag Registration

- o Application Protocol Tag: data.tcp
- o Intended Usage: See Section 6
- o Security Considerations: See Section 11
- o Contact Information: <one of the authors>

12.4. IPv4 Anycast

IANA has assigned a single IPv4 address from the 192.0.0.0/24 prefix and registered it in the "IANA IPv4 Special-Purpose Address Registry" [RFC6890].

Attribute	Value
Address Block Name	TBA Distributed-Denial-of-Service Open Threat Signaling (DOTS) Anycast
RFC	<this document>
Allocation Date	<date of approval of this document>
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Global	True
Reserved-by-Protocol	False

12.5. IPv6 Anycast

IANA has assigned a single IPv6 address from the 2001:0000::/23 prefix and registered it in the "IANA IPv6 Special-Purpose Address Registry" [RFC6890].

Attribute	Value
Address Block Name	TBA Distributed-Denial-of-Service Open Threat Signaling (DOTS) Anycast
RFC	<this document>
Allocation Date	<date of approval of this document>
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Global	True
Reserved-by-Protocol	False

13. Acknowledgements

Thanks to Brian Carpenter for the review of the BRSKI text.

Many thanks to Russ White for the review, comments, and text contribution.

14. References

14.1. Normative References

- [I-D.ietf-dots-architecture] Mortensen, A., Andreasen, F., K, R., christopher_gray3@cable.comcast.com, c., Compton, R., and N. Teague, "Distributed-Denial-of-Service Open Threat Signaling (DOTS) Architecture", draft-ietf-dots-architecture-07 (work in progress), September 2018.
- [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>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <<https://www.rfc-editor.org/info/rfc2132>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", RFC 3396, DOI 10.17487/RFC3396, November 2002, <<https://www.rfc-editor.org/info/rfc3396>>.
- [RFC3958] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", RFC 3958, DOI 10.17487/RFC3958, January 2005, <<https://www.rfc-editor.org/info/rfc3958>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, DOI 10.17487/RFC5986, September 2010, <<https://www.rfc-editor.org/info/rfc5986>>.

- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <<https://www.rfc-editor.org/info/rfc6890>>.
- [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>>.

14.2. Informative References

- [I-D.boucadair-dots-multihoming]
Boucadair, M. and R. K, "Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS)", draft-boucadair-dots-multihoming-03 (work in progress), April 2018.
- [I-D.ietf-anima-bootstrapping-keyinfra]
Pritikin, M., Richardson, M., Behringer, M., Bjarnason, S., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-16 (work in progress), June 2018.
- [I-D.ietf-dots-signal-channel]
K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-25 (work in progress), September 2018.
- [I-D.ietf-dots-use-cases]
Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-16 (work in progress), July 2018.
- [RFC2136] Vixie, P., Ed., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, DOI 10.17487/RFC2136, April 1997, <<https://www.rfc-editor.org/info/rfc2136>>.

- [RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, DOI 10.17487/RFC3007, November 2000, <<https://www.rfc-editor.org/info/rfc3007>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [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>>.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, DOI 10.17487/RFC4786, December 2006, <<https://www.rfc-editor.org/info/rfc4786>>.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<https://www.rfc-editor.org/info/rfc6125>>.
- [RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.
- [RFC7094] McPherson, D., Oran, D., Thaler, D., and E. Osterweil, "Architectural Considerations of IP Anycast", RFC 7094, DOI 10.17487/RFC7094, January 2014, <<https://www.rfc-editor.org/info/rfc7094>>.

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

Prashanth Patil
Cisco Systems, Inc.

Email: praspatti@cisco.com

DOTS
Internet-Draft
Intended status: Informational
Expires: September 10, 2019

M. Chen
Li. Su
CMCC
March 9, 2019

Using attack bandwidth in signal channel
draft-chen-dots-attack-bandwidth-expansion-00

Abstract

This document describes a DDoS Mitigation Request parameter used in the Signal Channel request, as an expansion of the signal channel for mitigating DDoS attack accurately with target-bandwidth. The proposed parameter will help to choose the mitigation method, to be blackhole directly or to be drained for clean.

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 September 10, 2019.

Copyright Notice

Copyright (c) 2019 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. Terminology	3
3. Mitigation Use Case	3
4. Request Mitigation expansion	5
5. Security Considerations	7
6. IANA Considerations	7
7. Acknowledgement	7
8. References	7
8.1. Normative References	7
8.2. Informative References	7
Authors' Addresses	8

1. Introduction

Distributed Denial of Service (DDoS) is a type of resource-consuming attack, which exploits a large number of attack resources and uses standard protocols to attack target objects. DDoS attacks consume a large amount of target object network resources or server resources (including computing power, storage capacity, etc.) of the target object, so that the target object cannot provide network services normally. At present, DDoS attack is one of the most powerful and indefensible attacks on the Internet, and due to the extensive use of mobile devices and IoT devices in recent years, it is easier for DDoS attackers to attack with real attack sources (broilers).

Volume based distributed denial-of-service attack bring huge amount of attack traffic on the link, and the peaks keep hitting new highs, the economic loss that causes is bigger also. For the service providers to immediately protect their network services from DDoS attacks, DDoS mitigation needs to be automated.

DDoS Open Threat Signaling (DOTS) is a protocol to standardize real-time signaling, threat-handling requests[I-D.ietf-dots-signal-channel], when attack target is under attack, dots client send mitigation request to dots server for help, If the mitigation request contains enough messages of the attack, then the mitigator can respond very effectively.

Currently, there are two selections to deal with ddos attacks on the link, one is blackhole, the other is flow clean. Blackhole means that all packets send to the attack target will be discarded by routers on the path, this way can instantly reduce the link load, Other managed services on this link will not be affected, but for the attack target all the normal business messages will be severely damaged, for example, if the attack target provide News and

information services and under ddos attack, all users will be inaccessible if the attack target choose blackhole for mitigation.

Flow clean means that all the flow will be drained by routers to clean center, the clean center will recognize the attack flow from normal business traffic, then reinjects normal business traffic to network link by routers after the operation of attack flow discard, in this way the attack target will not be effected.

This document describes attack-bandwidth, as a parameter expansion used in the mitigation request. attack-bandwidth means the amount of traffic under attack, this parameter can effectively reflect the degree of an attack, it will be more convenient for mitigator to choose the method for disposition when carry target-bandwidth in the mitigation request.

2. Terminology

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 [RFC2119]

The readers should be familiar with the terms defined in [I-D.ietf-dots-requirements] [I-D.ietf-dots-use-cases]

The terminology related to YANG data modules is defined in [RFC7950]

In addition, this document uses the terms defined below:

Attack-bandwidth: the amount of traffic under attack, it is usually expressed numerically.

Flow clean: one selection of Attack traffic deposition, the operation contains recognize, discard and reimage.

3. Mitigation Use Case

when attack target is under attack, it has to make corresponding disposal, there are two options for disposal, one is blackhole directly, in this way all the attack flow will be discarded by router upper path of attack target, this means that the attack target will not receive any traffic during the attack, all the traffic forwards attack target will be discarded, this has a huge impact on the work environment, especially the host that provide external service.

The other way of the disposition is to drainage all the traffic flow to clean center from router, then the clean center will use pattern

matching or any other method to find out the attack traffic flow to discard, finally, clean center reinage the normal business traffic back to attack target by upper router, the whole process above is defined as flow clean(Figure 1).




Figure 1: diagram of DDoS Mitigation usecase

Generally, the bandwidth of the link 1 must be larger than link 2 and link 3, and the clean ability of clean center limited to hardware resources. An example of link situation is as below(Figure 2):

figure tag	bandwidth/ capability
link 1	100Gb
link 2	50Gb
link 3	10Gb
clean center	80Gb

Figure 2: an example of link bandwidth

The Figure2 is a scenario of the link bandwidth, when a ddos attack is ongoing, if the link 1 bandwidth is completely jammed, the best way to mitigate the attack is to discard all the attack flow; if the amount of the traffic flow is lower than the remainder cleaning ability, the most suitable deposition is to drainage all the attack flow to clean center.

Therefore, it is an obvious requirement in the current network environment. In the architecture of DOTS, Dots client send mitigation request to dots server, the parameters in the mitigation request contains some message of attack target, but there have not

any messages of attack, if add attack-bandwidth to mitigation request as an expansion, it will be more effective and convenient for the disposition of mitigator.

4. Request Mitigation expansion

When a DOTS client requires mitigation for some reason, the DOTS client uses the CoAP PUT method to send a mitigation request to its DOTS server(s). If a DOTS client is entitled to solicit the DOTS service, the DOTS server enables mitigation on behalf of the DOTS client by communicating the DOTS client's request to a mitigator (which may be colocated with the DOTS server) and relaying the feedback of the thus-selected mitigator to the requesting DOTS client.

DOTS clients use the PUT method to request mitigation from a DOTS server. During active mitigation, DOTS clients may use PUT requests to carry mitigation efficacy updates to the DOTS server.

The new parameter in the CBOR body (Figure 3) is described below:

```

Content-Format: "application/dots+cbor"
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "string"
        ],
        "target-port-range": [
          {
            "lower-port": number,
            "upper-port": number
          }
        ],
        "target-protocol": [
          number
        ],
        "target-fqdn": [
          "string"
        ],
        "attack-bandwidth": [
          "string"
        ],
        "target-uri": [
          "string"
        ],
        "alias-name": [
          "string"
        ],
        "lifetime": number,
        "trigger-mitigation": true|false
      }
    ]
  }
}

```

Figure 3: PUT to Convey DOTS Mitigation Requests

attack-bandwidth: bandwidth occupied by an attack, The recommended format is numerical form, such as xxGb. Different attack has different attack bandwidth, numerical value directly reflects the urgency of the current attack. Serious attacks are treated with blackhole, Other cases use flow cleaning, attack-bandwidth is conducive to the selection of disposal mode.

This is an optional attribute.

The definition of the rest parameters are the same as the [I-D.ietf-dots-signal-channel]

5. Security Considerations

TBD

6. IANA Considerations

TBD

7. Acknowledgement

TBD

8. References

8.1. 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>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.

8.2. Informative References

- [I-D.ietf-dots-requirements] Mortensen, A., K, R., and R. Moskowitz, "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", draft-ietf-dots-requirements-20 (work in progress), February 2019.
- [I-D.ietf-dots-signal-channel] K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-30 (work in progress), March 2019.
- [I-D.ietf-dots-use-cases] Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-17 (work in progress), January 2019.

Authors' Addresses

Meiling Chen
CMCC
32, Xuanwumen West
BeiJing , BeiJing 100053
China

Email: chenmeiling@chinamobile.com

Li Su
CMCC
32, Xuanwumen West
BeiJing 100053
China

Email: suli@chinamobile.com

DOTS
Internet-Draft
Intended status: Informational
Expires: September 9, 2019

Y. Hayashi, Ed.
NTT
K. Nishizuka, Ed.
NTT Communications
M. Boucadair, Ed.
Orange
March 8, 2019

DDoS Mitigation Offload: A DOTS Applicability Use Case
draft-hayashi-dots-dms-offload-usecase-00

Abstract

This document describes the applicability of DOTS to a DDoS mitigation offload use case. This use case assumes that a DMS (DDoS Mitigation System) whose utilization rate is high sends its blocked traffic information to an orchestrator using DOTS protocols, then the orchestrator requests forwarding nodes such as routers to filter the traffic. Doing so enables service providers to mitigate DDoS attack traffic automatically while ensuring interoperability and distributed filter enforcement.

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 September 9, 2019.

Copyright Notice

Copyright (c) 2019 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. Terminology	3
3. The Problem	3
4. DOTS Applicability to DDoS Mitigation Offload Use Case	3
4.1. Component and Sequence Diagram	3
4.2. Case: DOTS Request via Out-of-band Link	5
4.3. Case: Mitigation Request via In-band Link	6
5. Security Considerations	7
6. IANA Considerations	7
7. Acknowledgement	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	8

1. Introduction

Volume-based distributed denial-of-service (DDoS) attacks such as DNS amplification attacks are critical threats to be handled by service providers. When such attacks occur, service providers have to mitigate them immediately to protect or recover their services.

Therefore, for the service providers to immediately protect their network services from DDoS attacks, DDoS mitigation needs to be automated. To automate DDoS attack mitigation, it is desirable that multi-vendor elements involved in DDoS attack detection and mitigation collaborate and support standard interfaces to communicate.

DDoS Open Threat Signaling (DOTS) is a set of protocols for real-time signaling, threat-handling requests, and data between the multi-vendor elements [I-D.ietf-dots-signal-channel] [I-D.ietf-dots-data-channel]. This document describes an automated DDoS Mitigation offload use case inherited from the DDoS orchestration use case [I-D.ietf-dots-use-cases], which ambitions to enable cost-effective DDoS Mitigation.

2. Terminology

The readers should be familiar with the terms defined in [I-D.ietf-dots-requirements] [I-D.ietf-dots-use-cases]

In addition, this document uses the terms defined below:

Mitigation offload: Getting rid of a DMS's mitigation action and assigning the action to another entity when the utilization rate of the DMS reaches a given threshold. How such threshold is set is deployment-specific.

Utilization rate: A scale to measure load of an entity such as link utilization rate or CPU utilization rate.

3. The Problem

In general, DDoS countermeasures are divided into detection and filtering, and detection is technically difficult. DDoS Mitigation System (DMS) can detect attack traffic based on the technology of their vendors, so service providers can increase DDoS countermeasure level by deploying the DMS in their network.

However, the number/capacity of DMS instances that can be deployed in a service providers network is limited due to equipment cost and dimensioning matters. Thus, DMS's utilization rate can reach its maximum capacity faster when the volume of DDoS attacks is enormous. When the rate reaches maximum capacity, the mitigation strategy needs to offload mitigation actions from the DMS to cost-effective forwarding nodes such as routers.

4. DOTS Applicability to DDoS Mitigation Offload Use Case

This section does not consider deployments where the network orchestrator and DMS are co-located.

4.1. Component and Sequence Diagram

Figures 1 and 2 show a component diagram and a sequence diagram of the use case, respectively.

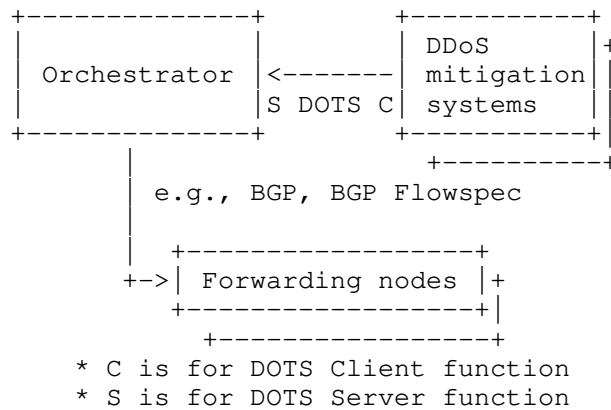


Figure 1: Component Diagram of DDoS Mitigation Offload Use Case

The component diagram shown in Figure 1 differs from that of DDoS Orchestration usecase in [I-D.ietf-dots-use-cases] in some respects. First, the DMS embeds a DOTS client to send DOTS requests to the orchestrator. Second, the orchestrator sends a request to underlying forwarding nodes to filter the attack traffic.

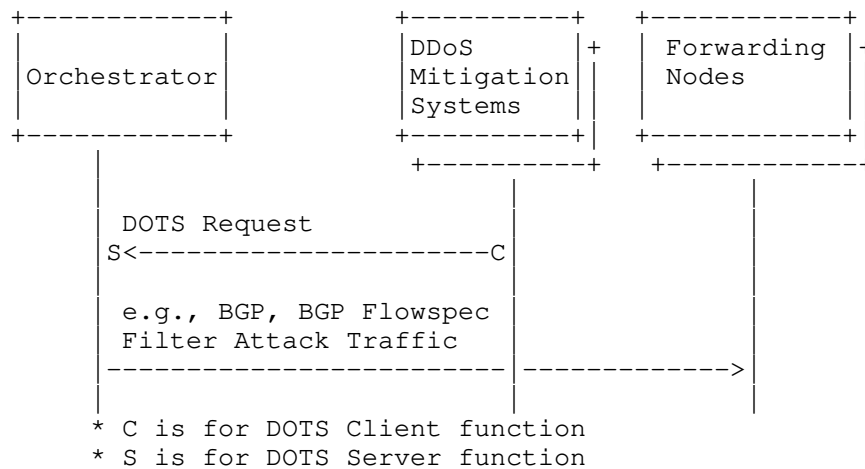


Figure 2: Sequence Diagram of DDoS Mitigation Offload Use Case

In this use case, it is assumed that volume based attack already hits a network and attack traffic is detected and blocked by a DMS in the network. When the volume-based attack becomes intense, DMS's

utilization rate can reach a certain threshold (e.g., maximum capacity). Then, the DMS sends a DOTS request as offload request to the orchestrator with the actions to enforce on the traffic. After that, the orchestrator requests the forwarding nodes to filter attack traffic by dissemination of flow specification rules protocols such as BGP Flowspec [RFC5575] on the basis of the blocked traffic information.

This use case is divided into two cases as discussed below. One is that the DMS sends DOTS requests to the orchestrator via out-of-band link, and the other one is that the DMS sends it via in-band link.

4.2. Case: DOTS Request via Out-of-band Link

In this case, the DMS sends a DOTS request to the orchestrator with information of blocked traffic information by the DMS via out-of-band link. The link is not congested when it is under volume attack-time, so DOTS data channel [I-D.ietf-dots-data-channel] is suitable because DOTS data channel has capability of conveying the drop-listed filtering rules (and other actions such as 'rate-limit'). The applicability of DOTS in such case is as follows:

- o The DMS generates a list of flow tuples (e.g., 5-tuples) which the DMS is blocking/rate-limiting and wants to offload.
- o The DMS creates ACEs for each elements of the list, setting "matches" as the flow tuple and "forwarding" in "actions" as "drop" (or other actions).
- o The DMS aggregates the ACEs under an ACL set, and the DMS sends the ACL to the orchestrator setting "activation-type" as "immediate".

Figure 3 shows a JSON example of ACL conveyed by DOTS data channel.

```

{
  "ietf-dots-data-channel:acls": {
    "acl": [
      {
        "name": "DMS_Offload_Usecase_ACL",
        "type": "ipv4-acl-type",
        "activation-type": "immediate",
        "aces": {
          "ace": [
            {
              "name": "DMS_Offload_Usecase_ACE_00",
              "matches": {
                "ipv4": {
                  "destination-ipv4-network": "192.0.2.2/32",
                  "source-ipv4-network": "203.0.113.2/32",
                  "protocol": 17
                },
                "udp": {
                  "source-port": {
                    "operator": "eq",
                    "port": 53
                  }
                }
              },
              "actions": {
                "forwarding": "drop"
              }
            }
          ]
        }
      ]
    ]
  }
}

```

Figure 3: JSON Example of ACL conveyed by DOTS data channel

4.3. Case: Mitigation Request via In-band Link

In this case, the DMS sends a mitigation request to the orchestrator with information of blocked traffic by the DMS via in-band channel. The link can be congested when it is under volume attack-time, so DOTS data channel can't be used to convey the drop-listed filtering rules as blocked traffic information [Interop].

The DOTS signal channel and [I-D.ietf-dots-signal-channel] and the source-* clauses defined in [I-D.reddy-dots-home-network] are used to communicate the policies to the orchestrator.

<<<An example will be included>>>>

5. Security Considerations

Security considerations discussed in [I-D.ietf-dots-data-channel] and [I-D.ietf-dots-signal-channel] are to be taken into account.

6. IANA Considerations

This document does not require any action from IANA.

7. Acknowledgement

Thanks to Tirumaleswar Reddy, Shunsuke Homma for the comments.
Thanks to Koichi Sakurada for demonstrating proof of concepts of this use case.

8. References

8.1. Normative References

[I-D.ietf-dots-data-channel]

Boucadair, M. and R. K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", draft-ietf-dots-data-channel-27 (work in progress), February 2019.

[I-D.ietf-dots-requirements]

Mortensen, A., K, R., and R. Moskowitz, "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", draft-ietf-dots-requirements-20 (work in progress), February 2019.

[I-D.ietf-dots-signal-channel]

K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-30 (work in progress), March 2019.

[I-D.ietf-dots-use-cases]

Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-17 (work in progress), January 2019.

8.2. Informative References

- [I-D.nishizuka-dots-signal-control-filtering]
Nishizuka, K., Boucadair, M., K, R., and T. Nagata,
"Controlling Filtering Rules Using DOTS Signal Channel",
draft-nishizuka-dots-signal-control-filtering-04 (work in
progress), February 2019.
- [I-D.reddy-dots-home-network]
K, R., Harsha, J., Boucadair, M., and J. Shallow, "Denial-
of-Service Open Threat Signaling (DOTS) Signal Channel
Call Home", draft-reddy-dots-home-network-03 (work in
progress), December 2018.
- [Interop] Nishizuka, K., Shallow, J., and L. Xia , "DOTS Interop
test report, IETF 103 Hackathon", November 2018,
<[https://datatracker.ietf.org/meeting/103/materials/
slides-103-dots-interop-report-from-ietf-103-hackathon-
00](https://datatracker.ietf.org/meeting/103/materials/slides-103-dots-interop-report-from-ietf-103-hackathon-00)>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A
Border Gateway Protocol 4 (BGP-4)", RFC 4271,
DOI 10.17487/RFC4271, January 2006,
<<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J.,
and D. McPherson, "Dissemination of Flow Specification
Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009,
<<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object
Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049,
October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.

Authors' Addresses

Yuhei Hayashi (editor)
NTT
3-9-11, Midori-cho
Musashino-shi, Tokyo 180-8585
Japan

Email: yuuei.hayashi@gmail.com, yuuei.hayashi.mr@hco.ntt.co.jp

Kaname Nishizuka (editor)
NTT Communications
GranPark 16F 3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan

Email: kaname@nttv6.jp

Mohamed Boucadair (editor)
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 25, 2019

M. Boucadair
Orange
T. Reddy
McAfee
January 21, 2019

Multi-homing Deployment Considerations for Distributed-Denial-of-Service
Open Threat Signaling (DOTS)
draft-ietf-dots-multihoming-01

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 July 25, 2019.

Copyright Notice

Copyright (c) 2019 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	4
4.1. Residential Single CPE	5
4.2. Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs	5
4.3. Multi-homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	6
4.4. Multi-homed Enterprise with the Same ISP	7
5. DOTS Deployment Considerations	7
5.1. Residential CPE	7
5.2. Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs	9
5.3. Multi-Homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	11
5.4. Multi-Homed Enterprise: Single ISP	12
6. Security Considerations	12
7. IANA Considerations	12
8. Acknowledgements	12
9. References	13
9.1. Normative References	13
9.2. Informative References	13
Authors' Addresses	14

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 seem to be under attack. To improve such situation, the IETF is specifying the DDoS Open Threat Signaling (DOTS) [I-D.ietf-dots-architecture] architecture, where a DOTS client can inform a DOTS server that the network is under a potential attack and that appropriate mitigation actions are required. Indeed, because the lack of a common method to coordinate a real-time response among involved actors and network domains jeopardizes the efficiency of DDoS attack mitigation actions, the DOTS protocol is meant to carry requests for DDoS attack mitigation, thereby reducing the impact of an attack and leading to more efficient responsive actions. [I-D.ietf-dots-use-cases] identifies a set of scenarios for DOTS; most of these scenarios involve a Customer Premises Equipment (CPE).

The high-level DOTS architecture is illustrated in Figure 1 ([I-D.ietf-dots-architecture]):




Figure 1: Basic DOTS Architecture

[I-D.ietf-dots-architecture] 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.

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:
 - * Send a DOTS mitigation request to an arbitrary DOTS server won't help 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. Sketch 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.

This document adopts the following methodology:

- o Identify and extract viable deployment candidates from [I-D.ietf-dots-use-cases].
- 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-signal-channel] and [I-D.ietf-dots-data-channel]; 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 [I-D.ietf-dots-architecture] and [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 sub-sections, 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. 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 (Customer Premises Equipment).
- o The CPE is connected to multiple provisioning domains (i.e., both fixed and mobile networks). Provisioning domain (PvD) is explained in [RFC7556].
- 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 [I-D.boucadair-dots-server-discovery]. 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 one 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.

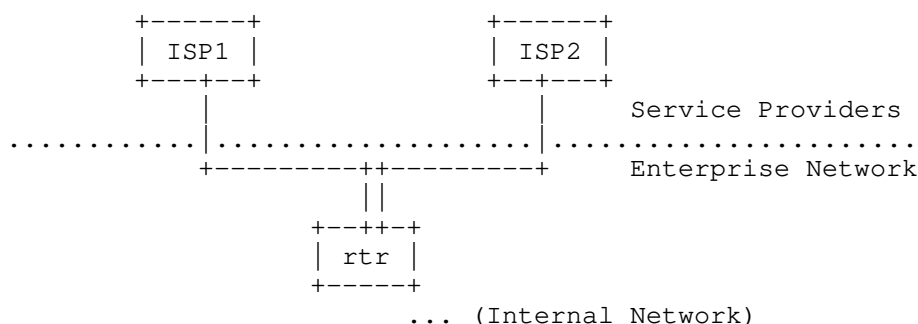


Figure 3: Multi-homed Enterprise Network (Single CPE connected to Multiple Networks)

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.

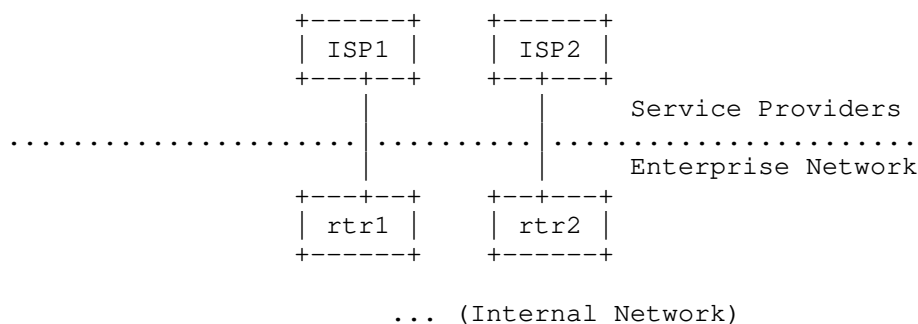


Figure 4: Multi-homed Enterprise Network (Multiple CPEs, Multiple ISPs)

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).

Editor's Note: The use of anycast addresses is to be consistently discussed.

5. DOTS 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 ([I-D.boucadair-dots-server-discovery]) using the DNS servers learned from the respective provisioning domain. The DOTS client 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 maintained with each of the DOTS servers because the mitigation scope of these servers is restricted. The DOTS client SHOULD use the certificate provisioned by a provisioning domain to authenticate itself to the DOTS server provided by the same provisioning domain.

When conveying a mitigation request to protect the attack target(s), the DOTS client among the DOTS servers available MUST select a DOTS server whose network has assigned the prefixes from which target prefixes and target IP addresses are derived. This implies that if no appropriate DOTS server is found, the DOTS client must not send the mitigation request to any 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 another domain 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. Nevertheless, if the DDoS attack is received from one single network, then only the DOTS server of that network must be contacted.

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 or 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 the same request to all its DOTS servers.




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 [I-D.boucadair-dots-server-discovery] 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 the mitigation request for all its PI addresses/prefixes to all the DOTS servers. The use of anycast addresses 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 arise 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 is provided with policies (e.g., prefix filter) that will trigger DOTS communications with the DOTS servers. Such policies will help the DOTS client to select the appropriate destination IP address.

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:

1. The same DOTS server for all network attachments.
2. 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 [I-D.ietf-dots-architecture].

TBD: In Home networks, if EST is used then how will the DOTS gateway (EST client) be provisioned with credentials for initial enrolment (see Section 2.2 in RFC 7030).

7. IANA Considerations

This document does not require any action from IANA.

8. Acknowledgements

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

Thanks to Kirill Kasavchenko for the comments.

9. References

9.1. Normative References

- [I-D.ietf-dots-architecture]
Mortensen, A., Andreasen, F., K, R., Teague, N., Compton, R., and c. christopher_gray3@cable.comcast.com, "Distributed-Denial-of-Service Open Threat Signaling (DOTS) Architecture", draft-ietf-dots-architecture-10 (work in progress), December 2018.
- [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>>.

9.2. Informative References

- [I-D.boucadair-dots-server-discovery]
Boucadair, M., K, R., and P. Patil, "Distributed-Denial-of-Service Open Threat Signaling (DOTS) Server Discovery", draft-boucadair-dots-server-discovery-05 (work in progress), October 2018.
- [I-D.ietf-dots-data-channel]
Boucadair, M., K, R., Nishizuka, K., Xia, L., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", draft-ietf-dots-data-channel-24 (work in progress), December 2018.
- [I-D.ietf-dots-signal-channel]
K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-26 (work in progress), December 2018.

[I-D.ietf-dots-use-cases]

Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-17 (work in progress), January 2019.

[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>>.

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

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.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised 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	10
5.3. Multi-Homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	12
5.4. Multi-Homed Enterprise: Single ISP	13
6. Security Considerations	14
7. IANA Considerations	14
8. Acknowledgements	14
9. References	15
9.1. Normative References	15
9.2. Informative References	15
Authors' Addresses	16

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 opsdireview, Mirja Kuhlewind for the tsvar 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

DOTS
Internet-Draft
Intended status: Standards Track
Expires: October 4, 2019

K. Nishizuka
NTT Communications
M. Boucadair
Orange
T. Reddy
McAfee
T. Nagata
Lepidum
April 2, 2019

Controlling Filtering Rules Using Distributed Denial-of-Service Open
Threat Signaling (DOTS) Signal Channel
draft-nishizuka-dots-signal-control-filtering-06

Abstract

This document specifies an extension to the DOTS signal channel so that DOTS clients can control their filtering rules when an attack mitigation is active.

Particularly, this extension allows a DOTS client to activate or deactivate existing filtering rules during a DDoS attack. The characterization of these filtering rules is supposed to be conveyed by a DOTS client during an idle time by means of the DOTS data channel protocol.

Editorial Note (To be removed by RFC Editor)

Please update these statements within the document with the RFC number to be assigned to this document:

- o "This version of this YANG module is part of RFC XXXX;"
- o "RFC XXXX: Controlling Filtering Rules Using Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel";
- o reference: RFC XXXX
- o [RFCXXXX]

Please update these statements with the RFC number to be assigned to the following documents:

- o "RFC SSSS: Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification" (used to be [I-D.ietf-dots-signal-channel])

- o "RFC DDDD: Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification" (used to be [I-D.ietf-dots-data-channel])

Please update the "revision" date of the YANG module.

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 October 4, 2019.

Copyright Notice

Copyright (c) 2019 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	3
1.1. The Problem	3
1.2. The Solution	4
2. Notational Conventions and Terminology	5
3. Controlling Filtering Rules of a DOTS Client	5
3.1. Binding the Data and Signal Channels	5
3.2. DOTS Signal Channel Extension	6
3.2.1. Parameters & Behaviors	6

3.2.2. DOTS Signal Filtering Control Module	8
3.2.2.1. Tree Structure	8
3.2.2.2. YANG Module	8
4. Sample Examples	11
4.1. Conflict Handling	11
4.2. On-Demand Activation of an Accept-List Filter	15
4.3. DOTS Servers/Mitigators Lacking Capacity	17
5. IANA Considerations	20
5.1. DOTS Signal Channel CBOR Mappings Registry	20
5.2. DOTS Signal Filtering Control YANG Module	21
6. Security Considerations	21
7. Acknowledgements	22
8. References	22
8.1. Normative References	22
8.2. Informative References	22
Authors' Addresses	23

1. Introduction

1.1. The Problem

The DOTS data channel protocol [I-D.ietf-dots-data-channel] is used for bulk data exchange between DOTS agents to improve the coordination of all the parties involved in the response to the DDoS attack. Filter management is one of its tasks which enables a DOTS client to retrieve the filtering capabilities of a DOTS server and to manage filtering rules. These Filtering rules are used for dropping or rate-limiting unwanted traffic, and permitting accept-listed traffic.

Unlike the DOTS signal channel [I-D.ietf-dots-signal-channel], the DOTS data channel is not expected to deal with attack conditions. As such, an issue that might be encountered in some deployments is when filters installed by means of DOTS data channel protocol may not function as expected during DDoS attacks or exacerbate an ongoing DDoS attack. The DOTS data channel cannot be used then to change these filters, which may complicate DDoS mitigation operations [Interop].

A typical case is a DOTS client which configures during 'idle' time (i.e., no mitigation is active) some filtering rules using DOTS data channel to permit traffic from accept-listed sources, but during a volumetric DDoS attack the DDoS mitigator identifies the source addresses/prefixes in the accept-listed filtering rules are attacking the target. For example, an attacker can spoof the IP addresses of accept-listed sources to generate attack traffic or the attacker can compromise the accept-listed sources and program them to launch a DDoS attack.

[I-D.ietf-dots-signal-channel] is designed so that the DDoS server notifies the conflict to the DOTS client (that is, 'conflict-cause' parameter set to 2 (Conflicts with an existing accept list)), but the DOTS client may not be able to withdraw the accept-list rules during the attack period due to the high-volume attack traffic saturating the inbound link. In other words, the DOTS client cannot use the DOTS data channel to withdraw the accept-list filters when the DDoS attack is in progress. This assumes that this DOTS client is the owner of the filtering rule.

1.2. The Solution

This specification addresses the problems discussed in Section 1.1 by adding the capability of managing filtering rules using the DOTS signal channel, which enables a DOTS client to request the activation or deactivation of filtering rules during a DDoS attack.

The DOTS signal channel protocol [I-D.ietf-dots-signal-channel] is designed to enable a DOTS client to contact a DOTS server for help even under severe network congestion conditions. Therefore, extending the DOTS signal channel protocol to manage the filtering rules during an attack will enhance the protection capability offered by DOTS protocols.

Note: The experiment at the IETF103 hackathon [Interop] showed that even when the incoming link is saturated by DDoS attack traffic, the DOTS client can signal mitigation requests using the DOTS signal channel over the saturated link.

Conflicts that are induced by filters installed by other DOTS clients of the same domain are not discussed in this specification.

Sample examples are provided in Section 4, in particular:

- o Section 4.1 illustrates how the filter control extension is used when conflicts with ACLs are detected by a DOTS server.
- o Section 4.2 shows how a DOTS client can instruct a DOTS server to safely forward some specific traffic in 'attack' time.
- o Section 4.3 shows how a DOTS client can react if DDoS traffic is still being forwarded to the DOTS client domain even if mitigation requests were sent to a DOTS server.

2. Notational Conventions and Terminology

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.

The reader should be familiar with the terms defined in [I-D.ietf-dots-requirements].

The meaning of the symbols in tree diagrams is defined in [RFC8340].

3. Controlling Filtering Rules of a DOTS Client

3.1. Binding the Data and Signal Channels

The filtering rules eventually managed using the DOTS signal channel are created a priori by the same DOTS client using the DOTS data channel. Managing conflicts with filters installed by other DOTS clients of the same domain is out of scope.

As discussed in Section 4.4.1 of [I-D.ietf-dots-signal-channel], a DOTS client must use the same 'cuid' for both the signal and data channels. This requirement is meant to facilitate binding DOTS channels used by the same DOTS client.

The DOTS signal and data channels from a DOTS client may or may not use the same DOTS server. Nevertheless, the scope of the mitigation request, alias, and filtering rules are not restricted to the DOTS server but to the DOTS server domain. To that aim, DOTS servers within a domain are assumed to have a mechanism to coordinate the mitigation requests, aliases, and filtering rules to coordinate their decisions for better mitigation operation efficiency. The exact details about such mechanism is out of scope of this document.

A filtering rule controlled by the DOTS signal channel is identified by its Access Control List (ACL) name (Section 7.2 of [I-D.ietf-dots-data-channel]). Note that an ACL name unambiguously identifies an ACL bound to a DOTS client, but the same name may be used by distinct DOTS clients.

The activation or deactivation of an ACL by the signal channel overrides the 'activation-type' (defined in Section 7.2 of [I-D.ietf-dots-data-channel]) a priori conveyed with the filtering rules using the DOTS data channel.

3.2. DOTS Signal Channel Extension

3.2.1. Parameters & Behaviors

This specification extends the mitigation request defined in Section 4.4.1 of [I-D.ietf-dots-signal-channel] to convey the intended control of the configured filtering rules. Concretely, the DOTS client conveys the following parameters in the CBOR body of a mitigation request:

acl-name: A name of an access list defined using the DOTS data channel (Section 7.2 of [I-D.ietf-dots-data-channel]).

As a reminder, an ACL is an ordered list of Access Control Entries (ACE). Each Access Control Entry has a list of match criteria and a list of actions [I-D.ietf-dots-data-channel]. The list of configured ACLs can be retrieved using the DOTS data channel during 'idle' time.

This is an optional attribute.

activation-type: Indicates the activation type of an ACL overriding the existing 'activation-type' installed by the DOTS client using the DOTS data channel.

This attribute can be set to 'deactivate', 'immediate', or 'activate-when-mitigating' defined [I-D.ietf-dots-data-channel].

Note that both 'immediate' and 'activate-when-mitigating' have an immediate effect when a mitigation request is being processed by the DOTS server.

If this attribute is not provided, the DOTS server MUST use 'activate-when-mitigating' as the default value.

This is an optional attribute.

The JSON/YANG mapping to CBOR for 'activation-type' is shown in Table 1.

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
activation-type	enumeration	0x0031 (TBD1)	0 unsigned	String

Table 1: JSON/YANG mapping to CBOR for 'activation-type'

A DOTS client may include `acl-*` attributes in a mitigation request having a new or an existing 'mid'. When `acl-*` attributes are to be included in a mitigation request with an existing 'mid', the DOTS client MUST repeat all the other parameters as sent in the original mitigation request (i.e., having that 'mid') apart from a possible change to the lifetime parameter value.

It is RECOMMENDED for a DOTS client to subscribe to asynchronous notifications of the attack mitigation, as detailed in Section 4.4.2.1 of [I-D.ietf-dots-signal-channel]. If not, the polling mechanism in Section 4.4.2.2 of [I-D.ietf-dots-signal-channel] has to be followed by the DOTS client.

A DOTS client MUST NOT use the filtering control over DOTS signal channel in 'idle' time; such requests MUST be discarded by the DOTS server with 4.00 (Bad Request). By default, ACL-related operations are achieved using the DOTS data channel [I-D.ietf-dots-data-channel] when no attack is ongoing.

A DOTS client relies on the information received from the DOTS server and/or local information to the DOTS client domain to trigger a filter control request. Only filters that are pertinent for an ongoing mitigation should be controlled by a DOTS client using the DOTS signal channel.

If the DOTS server does not find the ACL name conveyed in the mitigation request in its configuration data for this DOTS client, it MUST respond with a "4.04 (Not Found)" error response code.

If the DOTS server finds the ACL name for this DOTS client, and assuming the request passed the validation checks in [I-D.ietf-dots-signal-channel], the DOTS server MUST proceed with the 'activation-type' update. The update is immediately enforced by the DOTS server and will be maintained as the new activation type for the ACL name even after the termination of the mitigation request. In addition, the DOTS server MUST update the lifetime of the

corresponding ACL similar to the update when a refresh request is received using the DOTS data channel.

If, during an active mitigation, the 'activation-type' is changed at the DOTS server (e.g., as a result of an external action) for an ACL bound to a DOTS client, the DOTS server notifies that DOTS client with the change by including the corresponding acl-* parameters in an asynchronous notification (the DOTS client is observing the active mitigation) or in a response to a polling request (Section 4.4.2.2 of [I-D.ietf-dots-signal-channel]).

This specification does not require any modification to the efficacy update and the withdrawal procedures defined in [I-D.ietf-dots-signal-channel]. In particular, ACL-related clauses are not included in a PUT request used to send an efficacy update and DELETE requests.

3.2.2. DOTS Signal Filtering Control Module

3.2.2.1. Tree Structure

This document augments the "ietf-dots-signal-channel" DOTS signal YANG module defined in [I-D.ietf-dots-signal-channel] for managing filtering rules.

This document defines the YANG module "ietf-dots-signal-control", which has the following tree structure:

```
module: ietf-dots-signal-control
  augment /ietf-signal:dots-signal/ietf-signal:message-type
    /ietf-signal:mitigation-scope/ietf-signal:scope:
      +--rw acl-list* [acl-name] {control-filtering}?
        +--rw acl-name
        |   -> /ietf-data:dots-data/dots-client/acls/acl/name
        +--rw activation-type? ietf-data:activation-type
```

3.2.2.2. YANG Module

<CODE BEGINS> file "ietf-dots-signal-control@2019-04-01.yang"

```
module ietf-dots-signal-control {
  yang-version 1.1;
  namespace
    "urn:ietf:params:xml:ns:yang:ietf-dots-signal-control";
  prefix signal-control;

  import ietf-dots-signal-channel {
```

```
prefix ietf-signal;
reference
  "RFC SSSS: Distributed Denial-of-Service Open Threat
    Signaling (DOTS) Signal Channel Specification";
}
import ietf-dots-data-channel {
  prefix ietf-data;
  reference
    "RFC DDDD: Distributed Denial-of-Service Open Threat
      Signaling (DOTS) Data Channel Specification";
}

organization
  "IETF DDoS Open Threat Signaling (DOTS) Working Group";
contact
  "WG Web:    <https://datatracker.ietf.org/wg/dots/>
  WG List:    <mailto:dots@ietf.org>

  Author:     Konda, Tirumaleswar Reddy
              <mailto:TirumaleswarReddy_Konda@McAfee.com>

  Author:     Mohamed Boucadair
              <mailto:mohamed.boucadair@orange.com>

  Author:     Kaname Nishizuka
              <mailto:kaname@nttv6.jp>

  Author:     Takahiko Nagata
              <mailto:nagata@lepidum.co.jp>";

description
  "This module contains YANG definition for the signaling
  messages exchanged between a DOTS client and a DOTS server
  to control, by means of the DOTS signal channel, filtering
  rules configured using the DOTS data channel.

  Copyright (c) 2019 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX; see
  the RFC itself for full legal notices.";
```

```
revision 2019-04-01 {
  description
    "Initial revision.";
  reference
    "RFC XXXX: Controlling Filtering Rules Using Distributed
      Denial-of-Service Open Threat Signaling (DOTS)
      Signal Channel";
}

feature control-filtering {
  description
    "This feature means that the DOTS signal channel is able
      to manage the filtering rules created by the same DOTS
      client using the DOTS data channel.";
}

augment "/ietf-signal:dots-signal/ietf-signal:message-type"
  + "/ietf-signal:mitigation-scope/ietf-signal:scope" {
  if-feature control-filtering;

  description "ACL name and activation type.";

  list acl-list {
    key "acl-name";
    description
      "List of ACLs as defined in the DOTS data
        channel. These ACLs are uniquely defined by
        cuid and name.";
    leaf acl-name {
      type leafref {
        path "/ietf-data:dots-data/ietf-data:dots-client"
          + "/ietf-data:acls/ietf-data:acl/ietf-data:name";
      }
      description
        "Reference to the ACL name bound to a DOTS client.";
    }
    leaf activation-type {
      type ietf-data:activation-type;
      default "activate-when-mitigating";
      description
        "Set the activation type of an ACL.";
    }
  }
}
}
<CODE ENDS>
```

4. Sample Examples

This section provides sample examples to illustrate the behavior specified in Section 3.2.1. These examples are provided for illustration purposes; they should not be considered as deployment recommendations.

4.1. Conflict Handling

Let's consider a DOTS client which contacts its DOTS server during 'idle' time to install an accept-list allowing for UDP traffic issued from 2001:db8:1234::/48 with a destination port number 443 to be forwarded to 2001:db8:6401::2/127. It does so by sending, for example, a PUT request shown in Figure 1.

```

PUT /restconf/data/ietf-dots-data-channel:dots-data\
  /dots-client=paL8p4Zqo4SLv64TLPXrxA/acls\
  /acl=an-accept-list HTTP/1.1
Host: {host}:{port}
Content-Type: application/yang-data+json
{
  "ietf-dots-data-channel:acls": {
    "acl": [
      {
        "name": "an-accept-list",
        "type": "ipv6-acl-type",
        "activation-type": "activate-when-mitigating",
        "aces": {
          "ace": [
            {
              "name": "test-ace-ipv6-udp",
              "matches": {
                "ipv6": {
                  "destination-ipv6-network": "2001:db8:6401::2/127",
                  "source-ipv6-network": "2001:db8:1234::/48"
                },
                "udp": {
                  "destination-port": {
                    "operator": "eq",
                    "port": 443
                  }
                }
              },
              "actions": {
                "forwarding": "accept"
              }
            }
          ]
        }
      }
    ]
  }
}

```

Figure 1: DOTS Data Channel Request to Create a Filtering

Some time later, consider that a DDoS attack is detected by the DOTS client on 2001:db8:6401::2/127. Consequently, the DOTS client sends a mitigation request to its DOTS server as shown in Figure 2.

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "mitigate"
Uri-Path: "cuid=paL8p4Zqo4SLv64TLPXrxA"
Uri-Path: "mid=123"
Content-Format: "application/dots+cbor"
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "2001:db8:6401::2/127"
        ],
        "target-protocol": [
          17
        ],
        "lifetime": 3600
      }
    ]
  }
}
```

Figure 2: DOTS Signal Channel Mitigation Request

The DOTS server accepts immediately the request by replying with 2.01 (Created) (Figure 3).

```
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "mid": 123,
        "lifetime": 3600
      }
    ]
  }
}
```

Figure 3: Status Response (Message Body)

Assuming the DOTS client subscribed to asynchronous notifications, when the DOTS server concludes that some of the attack sources belong to 2001:db8:1234::/48, it sends a notification message with 'status' code set to '1 (Attack mitigation is in progress)' and 'conflict-cause' set to '2' (conflict-with-acceptlist) to the DOTS client to indicate that this mitigation request is in progress, but a conflict is detected.

Upon receipt of the notification message from the DOTS server, the DOTS client sends a PUT request to deactivate the "an-accept-list" ACL as shown in Figure 4.

The DOTS client can also decide to send a PUT request to deactivate the "an-accept-list" ACL, if suspect traffic is received from an accept-listed source (2001:db8:1234::/48). The structure of that PUT is the same as the one shown in Figure 4.

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "mitigate"
Uri-Path: "cuid=paL8p4Zqo4SLv64TLPXrxA"
Uri-Path: "mid=123"
Content-Format: "application/dots+cbor"
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "2001:db8:6401::2/127"
        ],
        "target-protocol": [
          17
        ],
        "acl-list": [
          {
            "acl-name": "an-accept-list",
            "activation-type": "deactivate"
          }
        ]
      }
    ]
  }
}
```

Figure 4: PUT for Deactivating a Conflicting Filter

Then, the DOTS server deactivates "an-accept-list" ACL and replies with 2.04 (Changed) response to the DOTS client to confirm the successful operation. The message body is similar to the one depicted in Figure 3.

Once the attack is mitigated, the DOTS client may use the data channel to retrieve its ACLs maintained by the DOTS server. As shown in Figure 5, the activation type is set to 'deactivate' as set by the

signal channel (Figure 4) instead of the type initially set using the data channel (Figure 1).

```
{
  "ietf-dots-data-channel:acls": {
    "acl": [
      {
        "name": "an-accept-list",
        "type": "ipv6-acl-type",
        "activation-type": "deactivate",
        "pending-lifetime": 10021,
        "aces": {
          "ace": [
            {
              "name": "test-ace-ipv6-udp",
              "matches": {
                "ipv6": {
                  "destination-ipv6-network": "2001:db8:6401::2/127",
                  "source-ipv6-network": "2001:db8:1234::/48"
                },
                "udp": {
                  "destination-port": {
                    "operator": "eq",
                    "port": 443
                  }
                }
              },
              "actions": {
                "forwarding": "accept"
              }
            }
          ]
        }
      }
    ]
  }
}
```

Figure 5: GET to Retrieve the Filtering (After Mitigation)

4.2. On-Demand Activation of an Accept-List Filter

Let's consider a DOTS client which contacts its DOTS server during 'idle' time to install an accept-list allowing for UDP traffic issued from 2001:db8:1234::/48 to be forwarded to 2001:db8:6401::2/127. It does so by sending, for example, a PUT request shown in Figure 6. The DOTS server installs this filter with a "deactivated" state.


```

PUT /restconf/data/ietf-dots-data-channel:dots-data\
  /dots-client=ioiuLoZqo4SLv64TLPXrxA/acls\
  /acl=my-accept-list HTTP/1.1
Host: {host}:{port}
Content-Type: application/yang-data+json
{
  "ietf-dots-data-channel:acls": {
    "acl": [
      {
        "name": "my-accept-list",
        "type": "ipv6-acl-type",
        "activation-type": "deactivate",
        "aces": {
          "ace": [
            {
              "name": "an-ace",
              "matches": {
                "ipv6": {
                  "destination-ipv6-network": "2001:db8:6401::2/127",
                  "source-ipv6-network": "2001:db8:1234::/48",
                  "protocol": 17
                }
              },
              "actions": {
                "forwarding": "accept"
              }
            }
          ]
        }
      }
    ]
  }
}

```

Figure 6: DOTS Data Channel Request to Create an Accept-List Filter

Sometime later, consider that a UDP DDoS attack is detected by the DOTS client on 2001:db8:6401::2/127 but the DOTS client wants to let the traffic from 2001:db8:1234::/48 to be accept-listed to the DOTS client domain. Consequently, the DOTS client sends a mitigation request to its DOTS server as shown in Figure 7.

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "mitigate"
Uri-Path: "cuid=ioiuLoZqo4SLv64TLPXrxA"
Uri-Path: "mid=4879"
Content-Format: "application/dots+cbor"
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "2001:db8:6401::2/127"
        ],
        "target-protocol": [
          17
        ],
        "acl-list": [
          {
            "acl-name": "my-accept-list",
            "activation-type": "immediate"
          }
        ]
      }
    ]
  }
}
```

Figure 7: DOTS Signal Channel Mitigation Request with a Filter Control

The DOTS server activates "my-accept-list" ACL and replies with 2.01 (Created) response to the DOTS client to confirm the successful operation.

4.3. DOTS Servers/Mitigators Lacking Capacity

This section describes a scenario in which a DOTS client activates a drop-list or a rate-limit filter during an attack.

Consider a DOTS client that contacts its DOTS server during 'idle' time to install an accept-list that rate-limits all (or a part thereof) traffic to be forwarded to 2001:db8:123::/48 as a last resort countermeasure whenever required. It does so by sending, for example, a PUT request shown in Figure 8. The DOTS server installs this filter with a "deactivated" state.

```
PUT /restconf/data/ietf-dots-data-channel:dots-data\  
  /dots-client=OopPisZqo4SLv64TLPXrxA/acls\  
  /acl=my-ratelimit-list HTTP/1.1  
Host: {host}:{port}  
Content-Type: application/yang-data+json  
{  
  "ietf-dots-data-channel:acls": {  
    "acl": [  
      {  
        "name": "my-ratelimit-list",  
        "type": "ipv6-acl-type",  
        "activation-type": "deactivate",  
        "aces": {  
          "ace": [  
            {  
              "name": "my-ace",  
              "matches": {  
                "ipv6": {  
                  "destination-ipv6-network": "2001:db8:123::/48"  
                }  
              },  
              "actions": {  
                "forwarding": "accept",  
                "rate-limit": "20.00"  
              }  
            }  
          ]  
        }  
      ]  
    }  
  }  
}
```

Figure 8: DOTS Data Channel Request to Create a Rate-Limit Filter

Consider now that a DDoS attack is detected by the DOTS client on 2001:db8:123::/48. Consequently, the DOTS client sends a mitigation request to its DOTS server (Figure 9).

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "mitigate"
Uri-Path: "cuid=OopPisZqo4SLv64TLPXrxA"
Uri-Path: "mid=85"
Content-Format: "application/dots+cbor"
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "2001:db8:123::/48"
        ],
        "lifetime": 3600
      }
    ]
  }
}
```

Figure 9: DOTS Signal Channel Mitigation Request to Activate a Rate-Limit Filter

For some reason (e.g., the DOTS server, or the mitigator, is lacking a capability or capacity), the DOTS client is still receiving the attack traffic which saturates available links. To soften the problem, the DOTS client decides to activate the filter that rate-limits the traffic destined to the DOTS client domain. To that aim, the DOTS client sends the mitigation request to its DOTS server shown in Figure 10.

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "mitigate"
Uri-Path: "cuid=OopPisZqo4SLv64TLPXrxA"
Uri-Path: "mid=85"
Content-Format: "application/dots+cbor"
{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "2001:db8:123::/48"
        ],
        "acl-list": [
          {
            "acl-name": "my-ratelimit-list",
            "activation-type": "activate"
          }
        ]
      }
    ]
  }
}
```

Figure 10: DOTS Signal Channel Mitigation Request to Activate a Rate-Limit Filter

Then, the DOTS server activates "my-ratelimit-list" ACL and replies with 2.04 (Changed) response to the DOTS client to confirm the successful operation.

5. IANA Considerations

5.1. DOTS Signal Channel CBOR Mappings Registry

This specification registers the 'activation-type' parameter in the IANA "DOTS Signal Channel CBOR Key Values" registry established by [I-D.ietf-dots-signal-channel].

The 'activation-type' is a comprehension-required parameter. The 'acl-list' and 'acl-name' parameters are defined as comprehension-required parameters in Table 6 in [I-D.ietf-dots-signal-channel]. Following the rules in [I-D.ietf-dots-signal-channel], if the DOTS server does not understand the 'acl-list' or 'acl-name' or 'activation-type' attributes, it responds with a "4.00 (Bad Request)" error response code.

- o Note to the RFC Editor: Please delete (TBD1) once the CBOR key is assigned from the (0x0001 - 0x3FFF) range.

Parameter Name	CBOR Key Value	CBOR Major Type	Change Controller	Specification Document (s)
activation-type	0x0031 (TBD1)	0	IESG	[RFCXXXX]

5.2. DOTS Signal Filtering Control YANG Module

This document requests IANA to register the following URI in the "IETF XML Registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-dots-signal-control
 Registrant Contact: The IESG.
 XML: N/A; the requested URI is an XML namespace.

This document requests IANA to register the following YANG module in the "YANG Module Names" registry [RFC7950].

Name: ietf-dots-signal-control
 Namespace: urn:ietf:params:xml:ns:yang:ietf-dots-signal-control
 Maintained by IANA: N
 Prefix: signal-control
 Reference: RFC XXXX

6. Security Considerations

The security considerations discussed in [I-D.ietf-dots-signal-channel] and [I-D.ietf-dots-data-channel] need to be taken into account.

A compromised DOTS client can use the filtering control capability to exacerbate an ongoing attack. Likewise, such compromised DOTS client may abstain from reacting to an ACL conflict notification received from the DOTS server during attacks. These are not new attack vectors, but variations of threats discussed in [I-D.ietf-dots-signal-channel] and [I-D.ietf-dots-data-channel]. DOTS operators should carefully monitor and audit DOTS agents to detect misbehavior and to deter misuse.

7. Acknowledgements

Thank you to Takahiko Nagata, Wei Pan, Xia Liang, and Jon Shollow for the comments.

8. References

8.1. Normative References

- [I-D.ietf-dots-data-channel]
Boucadair, M. and R. K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", draft-ietf-dots-data-channel-27 (work in progress), February 2019.
- [I-D.ietf-dots-signal-channel]
K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-30 (work in progress), March 2019.
- [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>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-dots-requirements]
Mortensen, A., K, R., and R. Moskowitz, "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", draft-ietf-dots-requirements-22 (work in progress), March 2019.

- [Interop] Nishizuka, K., Shallow, J., and L. Xia , "DOTS Interop test report, IETF 103 Hackathon", November 2018, <<https://datatracker.ietf.org/meeting/103/materials/slides-103-dots-interop-report-from-ietf-103-hackathon-00>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.

Authors' Addresses

Kaname Nishizuka
NTT Communications
GranPark 16F 3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan

Email: kaname@nttv6.jp

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: kondtir@gmail.com

Takahiko Nagata
Lepidum
Japan

Email: nagata@lepidum.co.jp

DOTS
Internet-Draft
Intended status: Standards Track
Expires: June 23, 2019

T. Reddy
J. Harsha
McAfee
M. Boucadair
Orange
J. Shallow
NCC Group
December 20, 2018

Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Call Home
draft-reddy-dots-home-network-03

Abstract

This document presents DOTS signal channel Call Home service, which enables a DOTS server to initiate a secure connection to a DOTS client, and to receive the attack traffic information from the DOTS client. The DOTS server in turn uses the attack traffic information to identify the compromised devices launching the outgoing DDOS attack and takes appropriate mitigation action.

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 June 23, 2019.

Copyright Notice

Copyright (c) 2018 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
1.1. The Problem	2
1.2. The Solution	4
1.3. Scope	4
2. Notational Conventions and Terminology	5
3. DOTS Signal Channel Call Home	5
3.1. Procedure	5
3.2. DOTS Signal Channel Extension	6
3.2.1. Mitigation Request	6
3.2.2. DOTS Signal Call Home YANG Module	9
4. IANA Considerations	12
4.1. DOTS Signal Channel Call Home UDP and TCP Port Number . .	12
4.2. DOTS Signal Channel CBOR Mappings Registry	12
4.3. New DOTS Conflict Cause	13
4.4. DOTS Signal Call Home YANG Module	13
5. Security Considerations	14
6. Privacy Considerations	14
7. Acknowledgements	15
8. References	15
8.1. Normative References	15
8.2. Informative References	16
Authors' Addresses	18

1. Introduction

1.1. The Problem

The DOTS signal channel protocol [I-D.ietf-dots-signal-channel] is used to carry information about a network resource or a network (or a part thereof) that is under a Distributed Denial of Service (DDoS) attack. Such information is sent by a DOTS client to one or multiple DOTS servers so that appropriate mitigation actions are undertaken on traffic deemed suspicious. Various use cases are discussed in [I-D.ietf-dots-use-cases].

IoT devices are becoming more and more prevalent in home networks, and with compute and memory becoming cheaper and cheaper, various types of IoT devices become available in the consumer market at affordable price. But on the downside, the main threat being most of these IoT devices are bought off-the-shelf and most manufacturers

haven't considered security in the product design. IoT devices deployed in home networks can be easily compromised, they do not have an easy mechanism to upgrade, and IoT manufactures may cease manufacture and/or discontinue patching vulnerabilities on IoT devices. However, these vulnerable and compromised devices will continue be used for a long period of time in the home, and the end-user does not know that IoT devices in his/her home are compromised. The compromised IoT devices are typically used for launching DDoS attacks on the victim while the owner/administrator of the home network is not aware about such misbehaviors. Similar to other DDoS attacks, the victim in this attack can be an application server, a host, a router, a firewall, or an entire network.

Nowadays, network devices in a home network offer network security, for instance, firewall/IPS service on a home router or gateway to protect the devices connected to the home network from external and internal attacks. Over the years several techniques have been identified to detect DDoS attacks, some of these techniques can be enabled on home network devices but most of them are used in the Internet Service Provider (ISP)'s network. The ISP offering DDoS mitigation service can detect outgoing DDoS attack traffic originating from its subscribers or the ISP may receive filtering rules (for example, using BGP flowspec [RFC5575]) from downstream service provider to filter, block, or rate-limit DDoS attack traffic originating from the ISP's subscribers to the downstream target.

Some of the DDoS attacks like spoofed RST or FIN packets, Slowloris, and TLS re-negotiation are difficult to detect on the home network devices without adversely affecting its performance. The reason is typically home routers have fast path to boost the throughput. For every new TCP/UDP flow, only the first few packets are punted through the slow path. Hence, it is not possible to detect various DDoS attacks in the slow path, since the attack payload is sent to the target server after the flow is switched to fast path. Deep packet inspection (DPI) of all the packets of a flow would be able to detect some of the attacks. However, a full-fledged DPI to detect these type of DDoS attacks is functionally or operationally not possible for all the devices attached to the home network owing to the memory and CPU limitations of the home routers. Further, for certain DDoS attacks the ability to distinguish legitimate traffic from attacker traffic on a per packet basis is complex. This complexity originates from the fact that the packet itself may look "legitimate" and no attack signature can be identified. The anomaly can be identified only after detailed statistical analysis.

The ISP on the other hand can detect the DDoS attack originating from a home network, but the ISP does not have a mechanism to detect which device in the home network is generating the DDoS attack traffic.

The primary reason being that devices in a IPv4 Home network are typically behind a NAT border. Even in case of a IPv6 Home network, although the ISP can identify the infected device in the Home network launching the DDoS traffic by tracking its unique IPv6 address, the infected device can easily change the IP address to evade remediation.

Existing approaches are still suffering from misused access network resources by abusing devices; the support of means for blocking such attacks close to the sources are missing. In particular, the DOTS signal protocol does not discuss cooperative DDoS mitigation between the home network and ISP to the suppress the outbound DDoS attack traffic originating from the home network.

1.2. The Solution

This specification addresses the problems discussed in Section 1.1 and presents DOTS signal channel Call Home extension, which enables the DOTS server to initiate a secure connection to the DOTS client, and the DOTS client then conveys the attack traffic information to the DOTS server.

In a typical deployment scenario, the DOTS server is enabled on a CPE, which is aligned with recent trends to enrich the CPE with advanced security features. Unlike classic DOTS deployments [I-D.ietf-dots-use-cases], such DOTS server maintains a single DOTS signal channel session for each DOTS-capable upstream provisioning domain [I-D.boucadair-dots-multihoming].

For instance, the DOTS server in the home network initiates the Call Home during peace time and then subsequently the DOTS client in the ISP environment can initiate a mitigation request whenever the ISP detects there is an attack from a compromised device in the DOTS server's domain.

The DOTS server uses the DDoS attack traffic information to identify the compromised device in its domain launching the DDoS attack, notifies the network administrator, and takes appropriate mitigation action. The mitigation action can be to quarantine the compromised device or block its traffic to the attack target until the mitigation request is withdrawn.

1.3. Scope

The aforementioned problems may be encountered in other deployments than those discussed Section 1.1. The solution proposed in this document can be used for those deployments to block DDoS attack traffic closer to the source(s) of the attack.

It is out of the scope of this document to identify an exhaustive list of such deployments.

2. Notational Conventions and Terminology

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.

The reader should be familiar with the terms defined in [I-D.ietf-dots-requirements].

3. DOTS Signal Channel Call Home

3.1. Procedure

DOTS signal channel Call Home preserves all but one of the DOTS client/server roles in the DOTS protocol stack, as compared to DOTS client-initiated DOTS signal channel protocol. The one and only role reversal that occurs are at the TCP/TLS or DTLS layers; that is, the DOTS server acts as a DTLS client and the DOTS client acts as a DTLS server or the DOTS server acts as a TCP/TLS client and the DOTS client acts as a TCP/TLS server. The DOTS server initiates TCP/TLS handshake or DTLS handshake to the DOTS client.

For example, a home network element (e.g., home router) co-located with a DOTS server (likely, a client-domain DOTS gateway) is the TCP/TLS server and DTLS server. However, when calling home, the DOTS server initially assumes the role of the TCP/TLS client and DTLS client, but the network element's role as a DOTS server remains the same. Further, existing certificate chains and mutual authentication mechanisms between the DOTS agents are unaffected by Call Home function. This Call Home function enables the DOTS server co-located with a network element (possibly behind NATs and firewalls) reachable by only the intended DOTS client and hence the DOTS server cannot be subjected to DDoS attacks. Other motivations for introducing Call Home are discussed in Section 1.1 of [RFC8071].

Figure 1 illustrates a sample Call Home flow exchange:

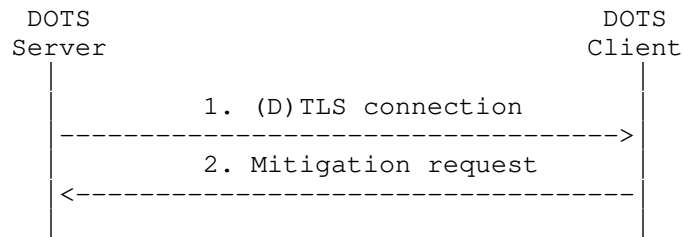


Figure 1: DOTS Signal Channel Call Home Sequence Diagram

The Call Home behavior is as follows:

1. If UDP transport is used, the DOTS server begins by initiating a DTLS connection to the DOTS client. The DOTS client MUST support accepting DTLS connection on the IANA-assigned port defined in Section 4.1, but MAY be configured to listen to a different port.

If TCP is used, the DOTS server begins by initiating a TCP connection to the DOTS client. The DOTS client MUST support accepting TCP connections on the IANA-assigned port defined in Section 4.1, but MAY be configured to listen to a different port. Using this TCP connection, the DOTS server initiates an TLS connection to the DOTS client.

The happy eyeballs mechanism explained in Section 4.3 of [I-D.ietf-dots-signal-channel] can be used for initiation of both TCP and UDP sessions.

2. Using this (D)TLS connection, the DOTS client requests, withdraws, or retrieves the status of mitigation requests.

3.2. DOTS Signal Channel Extension

3.2.1. Mitigation Request

This specification extends the mitigation request defined in [I-D.ietf-dots-signal-channel] to convey the attacker source prefixes and source port numbers. The DOTS client in the mitigation request conveys the following new parameters in the CBOR body of the mitigation request:

source-prefix: A list of attacker prefixes used to attack the target. Prefixes are represented using Classless Inter-Domain Routing (CIDR) notation [RFC4632].

As a reminder, the prefix length MUST be less than or equal to 32 (resp. 128) for IPv4 (resp. IPv6).

The prefix list MUST NOT include broadcast, loopback, or multicast addresses. These addresses are considered as invalid values. In addition, the DOTS client MUST validate that attacker prefixes are within the scope of the DOTS server's domain.

This is an optional attribute.

source-port-range: A list of port numbers used by the attack traffic flows.

A port range is defined by two bounds, a lower port number (lower-port) and an upper port number (upper-port). When only 'lower-port' is present, it represents a single port number.

For TCP, UDP, Stream Control Transmission Protocol (SCTP) [RFC4960], or Datagram Congestion Control Protocol (DCCP) [RFC4340], a range of ports can be, for example, 0-1023, 1024-65535, or 1024-49151.

This is an optional attribute.

source-icmp-type: A list of ICMP types used by the attack traffic flows. An ICMP type range is defined by two bounds, a lower ICMP type (lower-type) and an upper ICMP type (upper-type). When only 'lower-type' is present, it represents a single ICMP type.

This is an optional attribute.

The 'source-prefix' parameter is a mandatory attribute when the attack traffic information is signaled by the DOTS client in the call home scenario. The 'target-uri' or 'target-fqdn' parameters can be included in the mitigation request for diagnostic purpose to notify the DOTS server domain administrator, but SHOULD NOT be used to determine the target IP addresses. Note that 'target-prefix' becomes a mandatory attribute in the mitigation request signaling the attack information because 'target-uri' and 'target-fqdn' are optional attributes and 'alias-name' will not be conveyed in the mitigation request.

In order to help attack source identification by the DOTS server, the DOTS client SHOULD include in its mitigation request additional information such as 'source-port-range' or 'source-icmp-type-range'. The DOTS client MAY NOT include such information if 'source-prefix' conveys an IPv6 address/prefix.

If a Carrier Grade NAT (CGN, including NAT64) is located between the DOTS client domain and DOTS server domain, communicating an external IP address in a mitigation request is likely to be discarded by the

DOTS server because the external IP address is not visible locally to the DOTS server. The DOTS server is only aware of the internal IP addresses/prefixes bound to its domain. Thus, the DOTS client MUST NOT include the external IP address and/or port number identifying the suspect attack source, but MUST include the internal IP address and/or port number. To that aim, the DOTS client SHOULD rely on mechanisms, such as [I-D.ietf-opsawg-nat-yang] or [I-D.ietf-softwire-dslite-yang], to retrieve the internal IP address and port number which are mapped to an external IP address and port number.

If a MAP Border Relay [RFC7597] or lwAFTR [RFC7596] is enabled in the provider's domain to service its customers, the identification of an attack source bound to an IPv4 address/prefix MUST also rely on source port numbers because the same IPv4 address is assigned to multiple customers. The port information is required to unambiguously identify the source of an attack.

If a translator is enabled on the boundaries of the domain hosting the DOTS server (a CPE with NAT enabled, typically), the DOTS server uses the attack traffic information conveyed in a mitigation request to find the internal source IP address of the compromised device and blocks the traffic from the compromised device traffic to the attack target until the mitigation request is withdrawn. Doing so allows to isolate the suspicious device while avoiding to disturb other services.

The DOTS server domain administrator consent MAY be required to block the traffic from the compromised device to the attack target. An implementation MAY have a configuration knob to block the traffic from the compromised device to the attack target with or without DOTS server domain administrator consent. If the attack traffic is blocked, the DOTS server informs the DOTS client that the attack is being mitigated.

If the attack traffic information is identified by the DOTS server or the DOTS server domain administrator as legitimate traffic, the mitigation request is rejected, and 4.09 (Conflict) is returned to the DOTS client. The conflict-clause (defined in Section 4.4.1 of [I-D.ietf-dots-signal-channel]) indicates the cause of the conflict. The following new value is defined:

4: Mitigation request rejected. This code is returned by the DOTS server to indicate the attack traffic has been classified as legitimate traffic.

If the DOTS server is co-located with a home router, it can program the packet processor to punt all the traffic from the compromised

device to the target to slow path. The home router inspects the punted slow path traffic to detect and block the outgoing DDoS attack traffic or quarantine the device (e.g., using MAC level filtering) until it is remediated, and notifies the home administrator about the compromised device.

3.2.2. DOTS Signal Call Home YANG Module

3.2.2.1. Tree Structure

This document augments the "dots-signal-channel" DOTS signal YANG module defined in [I-D.ietf-dots-signal-channel] for signaling the attack traffic information. This document defines the YANG module "ietf-dots-signal-call-home", which has the following tree structure:

```
module: ietf-dots-signal-call-home
  augment /ietf-signal:dots-signal/ietf-signal:message-type
    /ietf-signal:mitigation-scope/ietf-signal:scope:
      +--rw source-prefix*          inet:ip-prefix {source-signaling}?
      +--rw source-port-range* [lower-port upper-port] {source-signaling}?
      |   +--rw lower-port          inet:port-number
      |   +--rw upper-port          inet:port-number
      +--rw source-icmp-type-range* [lower-type upper-type] {source-signaling}?
      |   +--rw lower-type          uint8
      |   +--rw upper-type          uint8
```

3.2.2.2. YANG Module

<CODE BEGINS> file "ietf-dots-signal-call-home@2018-09-28.yang"

```
module ietf-dots-signal-call-home {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-dots-signal-call-home";
  prefix signal-call-home;

  import ietf-inet-types {
    prefix inet;
    reference
      "Section 4 of RFC 6991";
  }
  import ietf-dots-signal-channel {
    prefix ietf-signal;
    reference
      "RFC YYYY: Distributed Denial-of-Service Open Threat
      Signaling (DOTS) Signal Channel Specification";
  }
}
```

organization

"IETF DDoS Open Threat Signaling (DOTS) Working Group";

contact

"WG Web: <<https://datatracker.ietf.org/wg/dots/>>
WG List: <<mailto:dots@ietf.org>>

Editor: Konda, Tirumaleswar Reddy
<mailto:TirumaleswarReddy_Konda@McAfee.com>;

Editor: Mohamed Boucadair
<<mailto:mohamed.boucadair@orange.com>>;

Editor: Jon Shallow
<<mailto:ietf-supjps@jpshallow.com>>;

description

"This module contains YANG definition for the signaling messages exchanged between a DOTS client and a DOTS server for the call home deployment scenario.

Copyright (c) 2018 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

revision 2018-09-28 {

description

"Initial revision.";

reference

"RFC XXXX: Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Call Home";

}

feature source-signaling {

description

"This feature means that source-related information can be supplied in mitigation requests.";

}

augment "/ietf-signal:dots-signal/ietf-signal:message-type/" +

```
    "ietf-signal:mitigation-scope/ietf-signal:scope" {
if-feature source-signaling;
description "Attacker source details";

leaf-list source-prefix {
    type inet:ip-prefix;
    description
        "IPv4 or IPv6 prefix identifying the attacker(s).";
}
list source-port-range {
    key "lower-port upper-port";
    description
        "Port range. When only lower-port is
        present, it represents a single port number.";
    leaf lower-port {
        type inet:port-number;
        mandatory true;
        description
            "Lower port number of the port range.";
    }
    leaf upper-port {
        type inet:port-number;
        must ". >= ../lower-port" {
            error-message
                "The upper port number must be greater than
                or equal to lower port number.";
        }
        description
            "Upper port number of the port range.";
    }
}
list source-icmp-type-range {
    key "lower-type upper-type";
    description
        "ICMP type range. When only lower-type is
        present, it represents a single ICMP type.";
    leaf lower-type {
        type uint8;
        mandatory true;
        description
            "Lower ICMP type of the ICMP type range.";
    }
    leaf upper-type {
        type uint8;
        must ". >= ../lower-type" {
            error-message
                "The upper ICMP type must be greater than
                or equal to lower ICMP type.";
        }
    }
}
```

```
    }
    description
      "Upper type of the ICMP type range.";
  }
}
}
}
<CODE ENDS>
```

4. IANA Considerations

4.1. DOTS Signal Channel Call Home UDP and TCP Port Number

IANA is requested to assign the port number TBD to the DOTS signal channel Call Home protocol for both UDP and TCP from the "Service Name and Transport Protocol Port Number Registry" available at: <https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>.

The assignment of port number 4647 is strongly suggested (DOTS signal channel uses port number 4646).

4.2. DOTS Signal Channel CBOR Mappings Registry

This specification registers the 'source-prefix' and 'source-port-range' parameters in the IANA "DOTS Signal Channel CBOR Mappings" registry established by [I-D.ietf-dots-signal-channel].

The 'source-prefix', 'source-port-range', and 'source-icmp-type-range' are comprehension-optional parameters.

- o Note to the RFC Editor: Please delete (TBD1)-(TBD5) once CBOR keys are assigned from the 0x8000 - 0xBFFF range.

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
source-prefix	leaf-list inet: ip-prefix	0x8000 (TBD1)	4 array	Array
source-port-range	list	0x8001 (TBD2)	3 text string 4 array	String Array
source-icmp-type-range	list	0x8002 (TBD3)	4 array	Array
lower-type	uint8	0x8003 (TBD4)	0 unsigned	Number
upper-type	uint8	0x8004 (TBD5)	0 unsigned	Number

4.3. New DOTS Conflict Cause

This document requests IANA to assign a new code from the "DOTS Conflict Cause Codes" registry:

Code	Label	Description	Reference
4	request-rejected	Mitigation request rejected. This code is returned by the DOTS server to indicate the attack traffic has been classified as legitimate traffic.	[RFCXXXX]

4.4. DOTS Signal Call Home YANG Module

This document requests IANA to register the following URI in the "IETF XML Registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-dots-signal-call-home
 Registrant Contact: The IESG.
 XML: N/A; the requested URI is an XML namespace.

This document requests IANA to register the following YANG module in the "YANG Module Names" registry [RFC7950].

```
name: ietf-signal-call-home
namespace: urn:ietf:params:xml:ns:yang:ietf-dots-signal-call-home
prefix: signal-call-home
reference: RFC XXXX
```

5. Security Considerations

This document deviates from standard DOTS signal channel usage by having the DOTS server initiate the TCP/TLS or DTLS connection. DOTS signal channel related security considerations discussed in Section 10 of [I-D.ietf-dots-signal-channel] MUST be considered. DOTS agents MUST authenticate each other using (D)TLS before a DOTS signal channel session is considered valid.

An attacker may launch a DoS attack on the DOTS client by having it perform computationally expensive operations, before deducing that the attacker doesn't possess a valid key. For instance, in TLS 1.3 [RFC8446], the ServerHello message contains a Key Share value based on an expensive asymmetric key operation for key establishment. Common precautions mitigating DoS attacks are recommended, such as temporarily blacklisting the source address after a set number of unsuccessful authentication attempts.

DOTS servers may not blindly trust mitigation requests from DOTS clients. For example, DOTS servers can use the attack flow information in a mitigation request to enable full-fledged packet inspection function to inspect all the traffic from the compromised to the target or to re-direct the traffic from the compromised device to the target to a DDoS mitigation system to scrub the suspicious traffic. DOTS servers can also seek the consent of DOTS server domain administrator to block the traffic from the compromised device to the target (see Section 3.2.1).

6. Privacy Considerations

Considerations discussed in [RFC6973] were taken into account to assess whether the DOTS Call Home extension introduces privacy threats.

Concretely, the protocol does not leak any new information that can be used to ease surveillance. In particular, the DOTS server is not required to share information that is local to its network (e.g., internal identifiers of an attack source) with the DOTS client.

The DOTS Call Home extension does not preclude the validation of mitigation requests received from a DOTS client. For example, a security service running on the CPE may require administrator's

consent before the CPE acts upon the mitigation request indicated by the DOTS client. How the consent is obtained is out of scope of this document.

Note that a DOTS server can seek for an administrator's consent, validate the request by inspecting the traffic, or proceed with both.

The DOTS Call Home extension is only advisory in nature. Concretely, the DOTS Call Home extension does not impose any action to be enforced within the home network; it is up to the DOTS server (and/or network administrator) to decide whether and which actions are required.

Moreover, the DOTS Call Home extension avoids misattribution by appropriately identifying the network to which a suspect attack source belongs to (e.g., address sharing issues discussed in Section 3.2.1).

Triggers to send a DOTS mitigation request to a DOTS server are deployment-specific. For example, a DOTS client may rely on the output of some DDoS detection systems deployed within the DOTS client's network to detect potential outbound DDoS attacks or on abuse claims received from remote victim networks. Such DDoS detection and mitigation techniques are not meant to track the activity of users, but to protect the Internet and avoid altering the IP reputation of the DOTS client's domain.

7. Acknowledgements

Thanks to Wei Pei, Xia Liang, Roman Danyliw, and Dan Wing for the comments.

8. References

8.1. Normative References

- [I-D.ietf-dots-signal-channel]
K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-25 (work in progress), September 2018.
- [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>>.

- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [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>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

8.2. Informative References

- [I-D.boucadair-dots-multihoming]
Boucadair, M. and R. K., "Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS)", draft-boucadair-dots-multihoming-04 (work in progress), October 2018.
- [I-D.ietf-dots-requirements]
Mortensen, A., Moskowitz, R., and R. K., "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", draft-ietf-dots-requirements-16 (work in progress), October 2018.
- [I-D.ietf-dots-use-cases]
Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-16 (work in progress), July 2018.
- [I-D.ietf-opsawg-nat-yang]
Boucadair, M., Sivakumar, S., Jacquenet, C., Vinapamula, S., and Q. Wu, "A YANG Module for Network Address Translation (NAT) and Network Prefix Translation (NPT)", draft-ietf-opsawg-nat-yang-17 (work in progress), September 2018.
- [I-D.ietf-softwire-dslite-yang]
Boucadair, M., Jacquenet, C., and S. Sivakumar, "A YANG Data Model for Dual-Stack Lite (DS-Lite)", draft-ietf-softwire-dslite-yang-17 (work in progress), May 2018.

- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, DOI 10.17487/RFC4340, March 2006, <<https://www.rfc-editor.org/info/rfc4340>>.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <<https://www.rfc-editor.org/info/rfc4632>>.
- [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>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <<https://www.rfc-editor.org/info/rfc7596>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<https://www.rfc-editor.org/info/rfc7597>>.
- [RFC8071] Watsen, K., "NETCONF Call Home and RESTCONF Call Home", RFC 8071, DOI 10.17487/RFC8071, February 2017, <<https://www.rfc-editor.org/info/rfc8071>>.

Authors' Addresses

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

Email: kondtir@gmail.com

Joshi Harsha
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: harsha_joshi@mcafee.com

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Jon Shallow
NCC Group
UK

Email: supjps-ietf@jpshallow.com

DOTS
Internet-Draft
Intended status: Standards Track
Expires: October 3, 2019

T. Reddy
McAfee
M. Boucadair
Orange
J. Shallow
April 01, 2019

Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal
Channel Call Home
draft-reddy-dots-home-network-04

Abstract

This document presents DOTS signal channel Call Home service, which enables a DOTS server to initiate a secure connection to a DOTS client, and to receive the attack traffic information from the DOTS client. The DOTS server in turn uses the attack traffic information to identify the compromised devices launching the outgoing DDoS attack and takes appropriate mitigation action.

The Call Home service is not specific to the home networks; the solution targets any deployment which requires to block DDoS attack traffic closer to the source(s) of a DDoS attack.

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 October 3, 2019.

Copyright Notice

Copyright (c) 2019 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
1.1. The Problem	2
1.2. The Solution	4
1.3. Scope	5
2. Notational Conventions and Terminology	5
3. DOTS Signal Channel Call Home	5
3.1. Procedure	5
3.2. DOTS Signal Channel Extension	6
3.2.1. Mitigation Request	6
3.2.2. DOTS Signal Call Home YANG Module	9
4. IANA Considerations	12
4.1. DOTS Signal Channel Call Home UDP and TCP Port Number	12
4.2. DOTS Signal Channel CBOR Mappings Registry	12
4.3. New DOTS Conflict Cause	13
4.4. DOTS Signal Call Home YANG Module	13
5. Security Considerations	14
6. Privacy Considerations	14
7. Contributors	15
8. Acknowledgements	15
9. References	16
9.1. Normative References	16
9.2. Informative References	16
Authors' Addresses	18

1. Introduction

1.1. The Problem

The DOTS signal channel protocol [I-D.ietf-dots-signal-channel] is used to carry information about a network resource or a network (or a part thereof) that is under a Distributed Denial of Service (DDoS) attack. Such information is sent by a DOTS client to one or multiple DOTS servers so that appropriate mitigation actions are undertaken on traffic deemed suspicious. Various use cases are discussed in [I-D.ietf-dots-use-cases].

IoT devices are becoming more and more prevalent in home networks, and with compute and memory becoming cheaper and cheaper, various types of IoT devices become available in the consumer market at affordable price. But on the downside, the main threat being most of these IoT devices are bought off-the-shelf and most manufacturers haven't considered security in the product design. IoT devices deployed in home networks can be easily compromised, they do not have an easy mechanism to upgrade, and IoT manufactures may cease manufacture and/or discontinue patching vulnerabilities on IoT devices. However, these vulnerable and compromised devices will continue be used for a long period of time in the home, and the end-user does not know that IoT devices in his/her home are compromised. The compromised IoT devices are typically used for launching DDoS attacks on the victim while the owner/administrator of the home network is not aware about such misbehaviors. Similar to other DDoS attacks, the victim in this attack can be an application server, a host, a router, a firewall, or an entire network.

Nowadays, network devices in a home network offer network security, for instance, firewall/IPS service on a home router or gateway to protect the devices connected to the home network from external and internal attacks. Over the years several techniques have been identified to detect DDoS attacks, some of these techniques can be enabled on home network devices but most of them are used in the Internet Service Provider (ISP)'s network. The ISP offering DDoS mitigation service can detect outgoing DDoS attack traffic originating from its subscribers or the ISP may receive filtering rules (for example, using BGP flowspec [RFC5575]) from downstream service provider to filter, block, or rate-limit DDoS attack traffic originating from the ISP's subscribers to the downstream target.

Some of the DDoS attacks like spoofed RST or FIN packets, Slowloris, and TLS re-negotiation are difficult to detect on the home network devices without adversely affecting its performance. The reason is typically home routers have fast path to boost the throughput. For every new TCP/UDP flow, only the first few packets are punted through the slow path. Hence, it is not possible to detect various DDoS attacks in the slow path, since the attack payload is sent to the target server after the flow is switched to fast path. Deep packet inspection (DPI) of all the packets of a flow would be able to detect some of the attacks. However, a full-fledged DPI to detect these type of DDoS attacks is functionally or operationally not possible for all the devices attached to the home network owing to the memory and CPU limitations of the home routers. Further, for certain DDoS attacks the ability to distinguish legitimate traffic from attacker traffic on a per packet basis is complex. This complexity originates from the fact that the packet itself may look "legitimate" and no

attack signature can be identified. The anomaly can be identified only after detailed statistical analysis.

The ISP on the other hand can detect the DDoS attack originating from a home network, but the ISP does not have a mechanism to detect which device in the home network is generating the DDoS attack traffic. The primary reason being that devices in a IPv4 Home network are typically behind a NAT border. Even in case of a IPv6 Home network, although the ISP can identify the infected device in the Home network launching the DDoS traffic by tracking its unique IPv6 address, the infected device can easily change the IP address to evade remediation.

Existing approaches are still suffering from misused access network resources by abusing devices; the support of means for blocking such attacks close to the sources are missing. In particular, the DOTS signal protocol does not discuss cooperative DDoS mitigation between the home network and ISP to suppress the outbound DDoS attack traffic originating from the home network.

1.2. The Solution

This specification addresses the problems discussed in Section 1.1 and presents DOTS signal channel Call Home extension, which enables the DOTS server to initiate a secure connection to the DOTS client, and the DOTS client then conveys the attack traffic information to the DOTS server.

In a typical deployment scenario, the DOTS server is enabled on a CPE, which is aligned with recent trends to enrich the CPE with advanced security features. Unlike classic DOTS deployments [I-D.ietf-dots-use-cases], such DOTS server maintains a single DOTS signal channel session for each DOTS-capable upstream provisioning domain [I-D.ietf-dots-multihoming].

For instance, the DOTS server in the home network initiates the Call Home in 'idle' time and then subsequently the DOTS client in the ISP environment can initiate a mitigation request whenever the ISP detects there is an attack from a compromised device in the DOTS server domain.

The DOTS server uses the DDoS attack traffic information to identify the compromised device in its domain launching the DDoS attack, notifies the network administrator, and takes appropriate mitigation action. The mitigation action can be to quarantine the compromised device or block its traffic to the attack target until the mitigation request is withdrawn.

1.3. Scope

The aforementioned problems may be encountered in other deployments than those discussed in Section 1.1. The solution specified in this document can be used for those deployments to block DDoS attack traffic closer to the source(s) of the attack.

It is out of the scope of this document to identify an exhaustive list of such deployments.

2. Notational Conventions and Terminology

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.

The reader should be familiar with the terms defined in [I-D.ietf-dots-requirements].

3. DOTS Signal Channel Call Home

3.1. Procedure

The DOTS signal channel Call Home extension preserves all but one of the DOTS client/server roles in the DOTS protocol stack, as compared to DOTS client-initiated DOTS signal channel protocol [I-D.ietf-dots-signal-channel]. The one and only role reversal that occurs are at the TCP/TLS or DTLS layers; that is, the DOTS server acts as a DTLS client and the DOTS client acts as a DTLS server or the DOTS server acts as a TCP/TLS client and the DOTS client acts as a TCP/TLS server. The DOTS server initiates TCP/TLS handshake or DTLS handshake to the DOTS client.

For example, a home network element (e.g., home router) co-located with a DOTS server (likely, a client-domain DOTS gateway) is the TCP/TLS server and DTLS server. However, when calling home, the DOTS server initially assumes the role of the TCP/TLS client and DTLS client, but the network element's role as a DOTS server remains the same. Furthermore, existing certificate chains and mutual authentication mechanisms between the DOTS agents are unaffected by the Call Home function. This Call Home function enables the DOTS server co-located with a network element (possibly behind NATs and firewalls) reachable by only the intended DOTS client and hence the DOTS server cannot be subjected to DDoS attacks. Other motivations for introducing the Call Home function are discussed in Section 1.1 of [RFC8071].

This document assumes that DOTS servers are provisioned with a way to know how to reach the upstream DOTS client(s), which could occur by a variety of means (e.g., [I-D.ietf-dots-server-discovery]). The specification of such means are out of scope of this document.

Figure 1 illustrates a sample Call Home flow exchange:

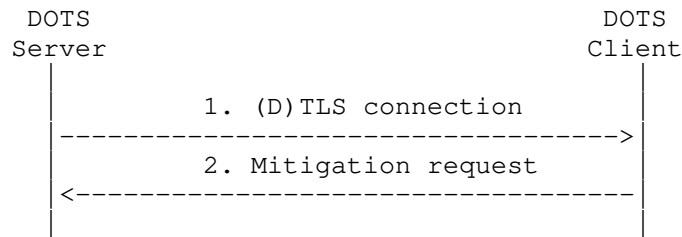


Figure 1: DOTS Signal Channel Call Home Sequence Diagram

The Call Home procedure is as follows:

1. If UDP transport is used, the DOTS server begins by initiating a DTLS connection to the DOTS client. The DOTS client **MUST** support accepting DTLS connection on the IANA-assigned port number defined in Section 4.1, but **MAY** be configured to listen to a different port number.

If TCP is used, the DOTS server begins by initiating a TCP connection to the DOTS client. The DOTS client **MUST** support accepting TCP connections on the IANA-assigned port number defined in Section 4.1, but **MAY** be configured to listen to a different port number. Using this TCP connection, the DOTS server initiates a TLS connection to the DOTS client.

The Happy Eyeballs mechanism explained in Section 4.3 of [I-D.ietf-dots-signal-channel] can be used for initiating (D)TLS connections.

2. Using this (D)TLS connection, the DOTS client may request, withdraw, or retrieve the status of mitigation requests.

3.2. DOTS Signal Channel Extension

3.2.1. Mitigation Request

This specification extends the mitigation request defined in [I-D.ietf-dots-signal-channel] to convey the attacker source prefixes and source port numbers. The DOTS client conveys the following new parameters in the CBOR body of the mitigation request:

source-prefix: A list of attacker prefixes used to attack the target. Prefixes are represented using Classless Inter-Domain Routing (CIDR) notation [RFC4632].

As a reminder, the prefix length MUST be less than or equal to 32 (resp. 128) for IPv4 (resp. IPv6).

The prefix list MUST NOT include broadcast, loopback, or multicast addresses. These addresses are considered as invalid values. In addition, the DOTS client MUST validate that attacker prefixes are within the scope of the DOTS server domain.

This is an optional attribute.

source-port-range: A list of port numbers used by the attack traffic flows.

A port range is defined by two bounds, a lower port number (lower-port) and an upper port number (upper-port). When only 'lower-port' is present, it represents a single port number.

For TCP, UDP, Stream Control Transmission Protocol (SCTP) [RFC4960], or Datagram Congestion Control Protocol (DCCP) [RFC4340], a range of ports can be, for example, 0-1023, 1024-65535, or 1024-49151.

This is an optional attribute.

source-icmp-type: A list of ICMP types used by the attack traffic flows. An ICMP type range is defined by two bounds, a lower ICMP type (lower-type) and an upper ICMP type (upper-type). When only 'lower-type' is present, it represents a single ICMP type.

This is an optional attribute.

The 'source-prefix' parameter is a mandatory attribute when the attack traffic information is signaled by a DOTS client in the Call Home scenario. The 'target-uri' or 'target-fqdn' parameters can be included in a mitigation request for diagnostic purposes to notify the DOTS server domain administrator, but SHOULD NOT be used to determine the target IP addresses. Note that 'target-prefix' becomes a mandatory attribute in the mitigation request signaling the attack information because 'target-uri' and 'target-fqdn' are optional attributes and 'alias-name' will not be conveyed in a mitigation request.

In order to help attack source identification by a DOTS server, the DOTS client SHOULD include in its mitigation request additional

information such as 'source-port-range' or 'source-icmp-type-range'. The DOTS client MAY NOT include such information if 'source-prefix' conveys an IPv6 address/prefix.

If a Carrier Grade NAT (CGN, including NAT64) is located between the DOTS client domain and DOTS server domain, communicating an external IP address in a mitigation request is likely to be discarded by the DOTS server because the external IP address is not visible locally to the DOTS server. The DOTS server is only aware of the internal IP addresses/prefixes bound to its domain. Thus, the DOTS client MUST NOT include the external IP address and/or port number identifying the suspect attack source, but MUST include the internal IP address and/or port number. To that aim, the DOTS client SHOULD rely on mechanisms, such as [RFC8512] or [RFC8513], to retrieve the internal IP address and port number which are mapped to an external IP address and port number.

If a MAP Border Relay [RFC7597] or lwAFTR [RFC7596] is enabled in the provider's domain to service its customers, the identification of an attack source bound to an IPv4 address/prefix MUST also rely on source port numbers because the same IPv4 address is assigned to multiple customers. The port information is required to unambiguously identify the source of an attack.

If a translator is enabled on the boundaries of the domain hosting the DOTS server (a CPE with NAT enabled, typically), the DOTS server uses the attack traffic information conveyed in a mitigation request to find the internal source IP address of the compromised device and blocks the traffic from the compromised device traffic to the attack target until the mitigation request is withdrawn. Doing so allows to isolate the suspicious device while avoiding to disturb other services.

The DOTS server domain administrator consent MAY be required to block the traffic from the compromised device to the attack target. An implementation MAY have a configuration knob to block the traffic from the compromised device to the attack target with or without DOTS server domain administrator consent. If the attack traffic is blocked, the DOTS server informs the DOTS client that the attack is being mitigated.

If the attack traffic information is identified by the DOTS server or the DOTS server domain administrator as legitimate traffic, the mitigation request is rejected, and 4.09 (Conflict) is returned to the DOTS client. The conflict-clause (defined in Section 4.4.1 of [I-D.ietf-dots-signal-channel]) indicates the cause of the conflict. The following new value is defined:

- 4: Mitigation request rejected. This code is returned by the DOTS server to indicate the attack traffic has been classified as legitimate traffic.

If the DOTS server is co-located with a home router, it can program the packet processor to punt all the traffic from the compromised device to the target to slow path. The home router inspects the punted slow path traffic to detect and block the outgoing DDoS attack traffic or quarantine the device (e.g., using MAC level filtering) until it is remediated, and notifies the home administrator about the compromised device.

3.2.2. DOTS Signal Call Home YANG Module

3.2.2.1. Tree Structure

This document augments the "dots-signal-channel" DOTS signal YANG module defined in [I-D.ietf-dots-signal-channel] for signaling the attack traffic information. This document defines the YANG module "ietf-dots-call-home", which has the following tree structure:

```
module: ietf-dots-call-home
  augment /ietf-signal:dots-signal/ietf-signal:message-type
    /ietf-signal:mitigation-scope/ietf-signal:scope:
      +--rw source-prefix*      inet:ip-prefix {source-signaling}?
      +--rw source-port-range*
         | [lower-port upper-port] {source-signaling}?
         +--rw lower-port      inet:port-number
         +--rw upper-port      inet:port-number
      +--rw source-icmp-type-range*
         | [lower-type upper-type] {source-signaling}?
         +--rw lower-type      uint8
         +--rw upper-type      uint8
```

3.2.2.2. YANG Module

```
<CODE BEGINS> file "ietf-dots-call-home@2018-04-01.yang"

module ietf-dots-call-home {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-dots-call-home";
  prefix call-home;

  import ietf-inet-types {
    prefix inet;
    reference
      "Section 4 of RFC 6991";
```

```
}
import ietf-dots-signal-channel {
  prefix ietf-signal;
  reference
    "RFC YYYY: Distributed Denial-of-Service Open Threat
      Signaling (DOTS) Signal Channel Specification";
}

organization
  "IETF DDoS Open Threat Signaling (DOTS) Working Group";
contact
  "WG Web:    <https://datatracker.ietf.org/wg/dots/>
  WG List:    <mailto:dots@ietf.org>

  Editor:     Konda, Tirumaleswar Reddy
              <mailto:TirumaleswarReddy_Konda@McAfee.com>;

  Editor:     Mohamed Boucadair
              <mailto:mohamed.boucadair@orange.com>;

  Editor:     Jon Shallow
              <mailto:ietf-supjps@jpshallow.com>";

description
  "This module contains YANG definition for the signaling
  messages exchanged between a DOTS client and a DOTS server
  for the call home deployment scenario.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX; see
  the RFC itself for full legal notices.";

revision 2018-04-01 {
  description
    "Initial revision.";
  reference
    "RFC XXXX: Distributed Denial-of-Service Open Threat
      Signaling (DOTS) Signal Channel Call Home";
}
```

```
feature source-signaling {
  description
    "This feature means that source-related information
    can be supplied in mitigation requests.";
}

augment "/ietf-signal:dots-signal/ietf-signal:message-type/"
  + "ietf-signal:mitigation-scope/ietf-signal:scope" {
  if-feature source-signaling;
  description "Attacker source details";

  leaf-list source-prefix {
    type inet:ip-prefix;
    description
      "IPv4 or IPv6 prefix identifying the attacker(s).";
  }
  list source-port-range {
    key "lower-port upper-port";
    description
      "Port range. When only lower-port is
      present, it represents a single port number.";
    leaf lower-port {
      type inet:port-number;
      mandatory true;
      description
        "Lower port number of the port range.";
    }
    leaf upper-port {
      type inet:port-number;
      must ". >= ../lower-port" {
        error-message
          "The upper port number must be greater than
          or equal to lower port number.";
      }
      description
        "Upper port number of the port range.";
    }
  }
}
list source-icmp-type-range {
  key "lower-type upper-type";
  description
    "ICMP type range. When only lower-type is
    present, it represents a single ICMP type.";
  leaf lower-type {
    type uint8;
    mandatory true;
    description
      "Lower ICMP type of the ICMP type range.";
  }
}
```

```
    }  
    leaf upper-type {  
      type uint8;  
      must ". >= ../lower-type" {  
        error-message  
          "The upper ICMP type must be greater than  
          or equal to lower ICMP type.";  
      }  
      description  
        "Upper type of the ICMP type range.";  
    }  
  }  
}  
}  
}  
<CODE ENDS>
```

4. IANA Considerations

4.1. DOTS Signal Channel Call Home UDP and TCP Port Number

IANA is requested to assign the port number TBD to the DOTS signal channel Call Home protocol for both UDP and TCP from the "Service Name and Transport Protocol Port Number Registry" available at: <https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>.

The assignment of port number 4647 is strongly suggested (DOTS signal channel uses port number 4646).

4.2. DOTS Signal Channel CBOR Mappings Registry

This specification registers the 'source-prefix' and 'source-port-range' parameters in the IANA "DOTS Signal Channel CBOR Mappings" registry established by [I-D.ietf-dots-signal-channel].

The 'source-prefix', 'source-port-range', and 'source-icmp-type-range' are comprehension-optional parameters.

- o Note to the RFC Editor: Please delete (TBD1)-(TBD5) once CBOR keys are assigned from the 0x8000 - 0xBFFF range.

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
source-prefix	leaf-list inet: ip-prefix	0x8000 (TBD1)	4 array	Array
source-port-range	list	0x8001 (TBD2)	3 text string 4 array	String Array
source-icmp-type-range	list	0x8002 (TBD3)	4 array	Array
lower-type	uint8	0x8003 (TBD4)	0 unsigned	Number
upper-type	uint8	0x8004 (TBD5)	0 unsigned	Number

4.3. New DOTS Conflict Cause

This document requests IANA to assign a new code from the "DOTS Conflict Cause Codes" registry:

Code	Label	Description	Reference
4	request-rejected	Mitigation request rejected. This code is returned by the DOTS server to indicate the attack traffic has been classified as legitimate traffic.	[RFCXXXX]

4.4. DOTS Signal Call Home YANG Module

This document requests IANA to register the following URI in the "IETF XML Registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-dots-call-home
 Registrant Contact: The IESG.
 XML: N/A; the requested URI is an XML namespace.

This document requests IANA to register the following YANG module in the "YANG Module Names" registry [RFC7950].

Name: ietf-call-home
Namespace: urn:ietf:params:xml:ns:yang:ietf-dots-call-home
Maintained by IANA: N
Prefix: call-home
Reference: RFC XXXX

5. Security Considerations

This document deviates from classic DOTS signal channel usage by having the DOTS server initiate the TCP/TLS or DTLS connection. DOTS signal channel related security considerations discussed in Section 10 of [I-D.ietf-dots-signal-channel] MUST be considered. DOTS agents MUST authenticate each other using (D)TLS before a DOTS signal channel session is considered valid.

An attacker may launch a DoS attack on the DOTS client by having it perform computationally expensive operations, before deducing that the attacker doesn't possess a valid key. For instance, in TLS 1.3 [RFC8446], the ServerHello message contains a Key Share value based on an expensive asymmetric key operation for key establishment. Common precautions mitigating DoS attacks are recommended, such as temporarily blacklisting the source address after a set number of unsuccessful authentication attempts.

DOTS servers may not blindly trust mitigation requests from DOTS clients. For example, DOTS servers can use the attack flow information in a mitigation request to enable full-fledged packet inspection function to inspect all the traffic from the compromised to the target or to re-direct the traffic from the compromised device to the target to a DDoS mitigation system to scrub the suspicious traffic. DOTS servers can also seek the consent of DOTS server domain administrator to block the traffic from the compromised device to the target (see Section 3.2.1).

6. Privacy Considerations

The considerations discussed in [RFC6973] were taken into account to assess whether the DOTS Call Home extension introduces privacy threats.

Concretely, the protocol does not leak any new information that can be used to ease surveillance. In particular, the DOTS server is not required to share information that is local to its network (e.g., internal identifiers of an attack source) with the DOTS client.

The DOTS Call Home extension does not preclude the validation of mitigation requests received from a DOTS client. For example, a security service running on the CPE may require administrator's

consent before the CPE acts upon the mitigation request indicated by the DOTS client. How the consent is obtained is out of scope of this document.

Note that a DOTS server can seek for an administrator's consent, validate the request by inspecting the traffic, or proceed with both.

The DOTS Call Home extension is only advisory in nature. Concretely, the DOTS Call Home extension does not impose any action to be enforced within the home network; it is up to the DOTS server (and/or network administrator) to decide whether and which actions are required.

Moreover, the DOTS Call Home extension avoids misattribution by appropriately identifying the network to which a suspect attack source belongs to (e.g., address sharing issues discussed in Section 3.2.1).

Triggers to send a DOTS mitigation request to a DOTS server are deployment-specific. For example, a DOTS client may rely on the output of some DDoS detection systems deployed within the DOTS client's network to detect potential outbound DDoS attacks or on abuse claims received from remote victim networks. Such DDoS detection and mitigation techniques are not meant to track the activity of users, but to protect the Internet and avoid altering the IP reputation of the DOTS client's domain.

7. Contributors

The following individuals have contributed to this document:

Joshi Harsha
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: harsha_joshi@mcafee.com

8. Acknowledgements

Thanks to Wei Pei, Xia Liang, Roman Danyliw, Dan Wing, and Toema Gavrichenkov for the comments.

9. References

9.1. Normative References

- [I-D.ietf-dots-signal-channel]
K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-30 (work in progress), March 2019.
- [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>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [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>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

9.2. Informative References

- [I-D.ietf-dots-multihoming]
Boucadair, M. and R. K., "Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS)", draft-ietf-dots-multihoming-01 (work in progress), January 2019.
- [I-D.ietf-dots-requirements]
Mortensen, A., K, R., and R. Moskowitz, "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", draft-ietf-dots-requirements-22 (work in progress), March 2019.

- [I-D.ietf-dots-server-discovery]
Boucadair, M., K, R., and P. Patil, "Distributed-Denial-of-Service Open Threat Signaling (DOTS) Server Discovery", draft-ietf-dots-server-discovery-00 (work in progress), March 2019.
- [I-D.ietf-dots-use-cases]
Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-17 (work in progress), January 2019.
- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, DOI 10.17487/RFC4340, March 2006, <<https://www.rfc-editor.org/info/rfc4340>>.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <<https://www.rfc-editor.org/info/rfc4632>>.
- [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>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <<https://www.rfc-editor.org/info/rfc7596>>.

- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<https://www.rfc-editor.org/info/rfc7597>>.
- [RFC8071] Watsen, K., "NETCONF Call Home and RESTCONF Call Home", RFC 8071, DOI 10.17487/RFC8071, February 2017, <<https://www.rfc-editor.org/info/rfc8071>>.
- [RFC8512] Boucadair, M., Ed., Sivakumar, S., Jacquenet, C., Vinapamula, S., and Q. Wu, "A YANG Module for Network Address Translation (NAT) and Network Prefix Translation (NPT)", RFC 8512, DOI 10.17487/RFC8512, January 2019, <<https://www.rfc-editor.org/info/rfc8512>>.
- [RFC8513] Boucadair, M., Jacquenet, C., and S. Sivakumar, "A YANG Data Model for Dual-Stack Lite (DS-Lite)", RFC 8513, DOI 10.17487/RFC8513, January 2019, <<https://www.rfc-editor.org/info/rfc8513>>.

Authors' Addresses

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

Email: kondtir@gmail.com

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Jon Shallow
UK

Email: supjps-ietf@jpshallow.com