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

M. Boucadair
Orange
T. Reddy
McAfee

October 7, 2018

Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS) draft-boucadair-dots-multihoming-04

Abstract

This document discusses multi-homing considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS). The goal is to provide a set of 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 April 10, 2019.

Table of Contents

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; almost all these scenarios involve a CPE.

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

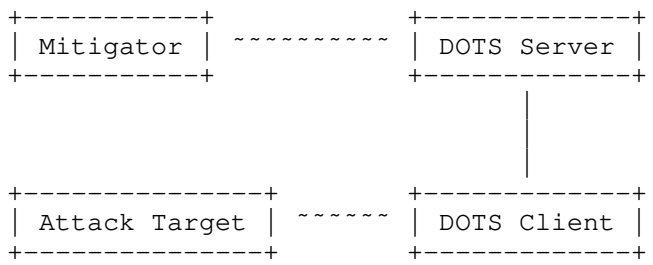

Figure 1: Basic DOTS Architecture

[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 signaling 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 taking 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 service 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.

To that aim, this document adopts the following methodology:

- Identify and extract viable deployment candidates from [I-D.ietf-dots-use-cases].
- 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)
- 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 multihomed 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 refers to both IPv4 and IPv6.

4. Multi-Homing Scenarios

This section briefly 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 hereafter.

4.1. Residential CPE

The scenario shown in Figure 2 is characterized as follows:

- The home network is connected to the Internet using one single CPE (Customer Premises Equipment).

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

- Each of these provisioning domains assign IP addresses/prefixes to the CPE. These addresses/prefixes are said to be Provider-Aggregatable (PA).

- The CPE is provided by each of these provisioning domains with 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].

- Because of ingress filtering, packets forwarded by the CPE to a given provisioning domain must be send with a source IP address that was assigned by that network [RFC8043].

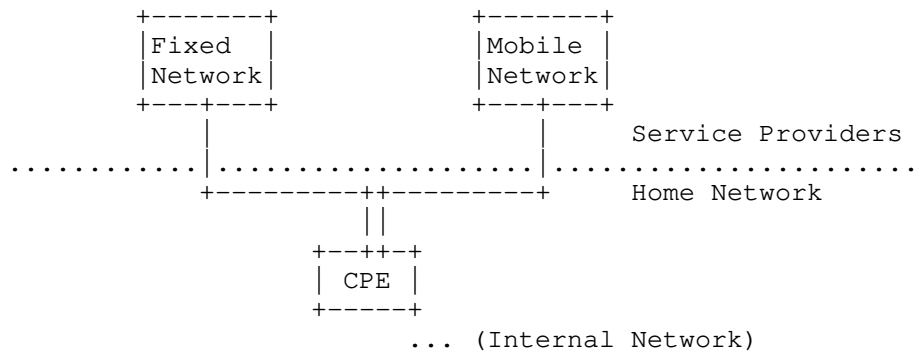

Figure 2: Typical Multi-homed Residential CPE

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 one single router.

- 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. Provider Independent (PI) addresses: The enterprise is the owner of the IP addresses/prefixes; the same address/prefix is then used for communication placed using any of the provisioning domains.

2. PA addresses/prefixes: each of provisioning domains assigns IP addresses/prefixes to the enterprise network.

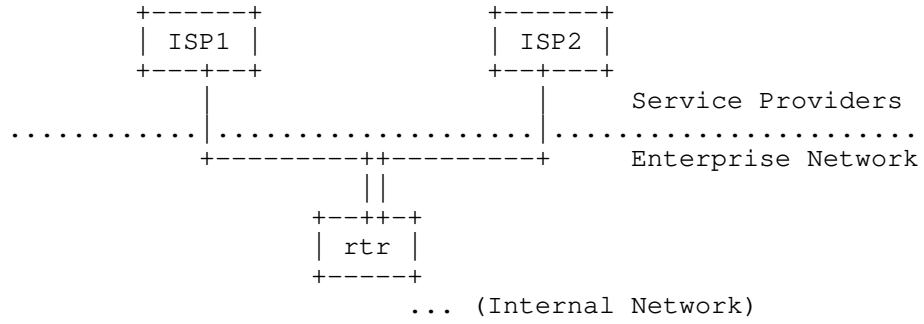

Figure 3: Multi-homed Enterprise Network (Single CPE connected to Multiple Networks)

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

This scenario is similar to the one in Section 4.2; the main difference is that dedicated routers are used to connect to each provisioning domain.

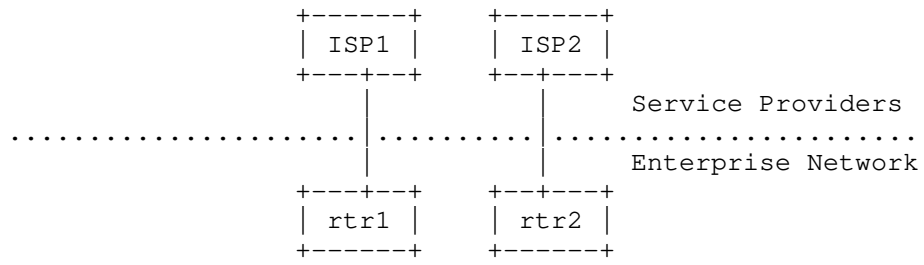

Figure 4: Multi-homed Enterprise Network (Multiple CPEs, Multiple ISPs)

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 provided by the same ISP (i.e., same provisioning domain).

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 sub-sections.

5.1. Residential CPE

Figure 5 depicts DOTS signaling sessions that are required to be established between a DOTS client (C) and DOTS servers (S1, S2) in the context of the scenario described in Section 4.1.

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 same 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 signaling 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. For example, mitigation request to protect target resources bound to a PA IP address/prefix cannot be honored by an provisioning domain other than the one that owns those addresses/prefixes. Consequently, Typically, if a CPE detects a DDoS attack on all its network attachments, it must contact both DOTS servers for mitigation. 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.

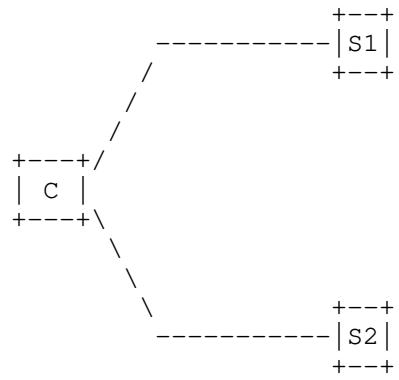

Figure 5: DOTS associations for a multihomed residential CPE

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 is enabled in the context of the scenario described in Section 4.2. This deployment is characterized as follows:

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

- A DOTS getaway is enabled to aggregate/relay the requests to upstream DOTS servers.

When PA addresses/prefixes are in used, the same considerations discussed in Section 5.1 are to be followed by the DOTS gateway to contact its DOTS server(s). The DOTS gateways can be reachable from DOTS client using a unicast or anycast address.

Nevertheless, when PI addresses/prefixes are assigned, the DOTS gateway MUST sent the same request to all its DOTS servers.

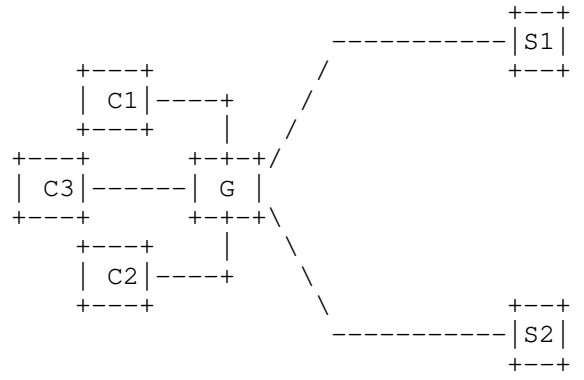

Figure 6: Multiple DOTS Clients, Single DOTS Gateway, Multiple DOTS Servers

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

- One of more DOTS clients are enabled in hosts located in the internal network. These DOTS client may use [I-D.boucadair-dots-server-discovery] to discover its DOTS server(s).

- These DOTS clients communicate directly with upstream DOTS servers.

If PI addresses/prefixes are in use, the DOTS client can send the mitigation request for all its PI addresses/prefixes to any one of the DOTS servers. The use of anycast addresses is NOT RECOMMENDED.

If PA addresses/prefxies are used, the same considerations discussed in Section 5.1 are to be followed by the DOTS clients. Because DOTS clients are not located on the CPE and multiple addreses/prefixes may not be assigned to the DOTS client (IPv4 context, typically), some complications arise to steer the traffic to the appropriate DOTS

server using the appropriate source IP address. These complications discussed in [RFC4116] are not specific to DOTS .

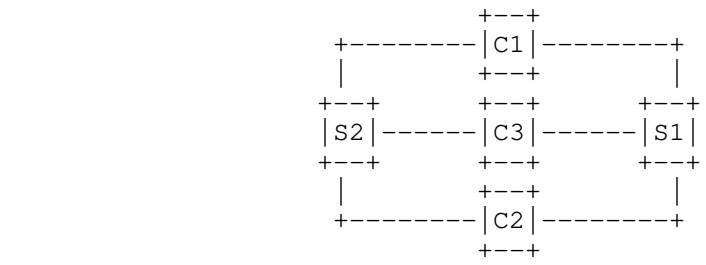

Figure 7: Multiple DOTS Clients, Multiple DOTS Servers

Another deployment approach is to enable many DOTS clients; each of them responsible to handle communication with a specific DOTS server (see Figure 8). Each DOTS client is provided with policies (e.g., prefix filter) that will trigger DOTS communications with the DOTS servers. 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.

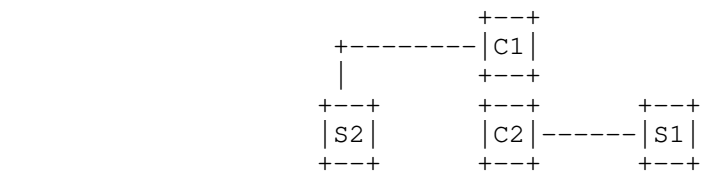

Figure 8: Single Homed DOTS Clients

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

The deployments depicted in Figure 7 and Figure 8 apply also for 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 DOTS gateway is enabled in each CPE (rtr1, rtr2).

o Each of these DOTS gateways communicate with the DOTS server of the provisioning domain.

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

When PA addresses/prefixes are used, but no filter rules are provided to DOTS clients, these later 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 or be rejected.

When PA addresses/prefixes are used, but specific filter rules are provided to DOTS clients using some means that are out of scope, these later MUST select the appropriate DOTS gateway to be contacted. The use of anycast is NOT RECOMMENDED to reach DOTS gateways.

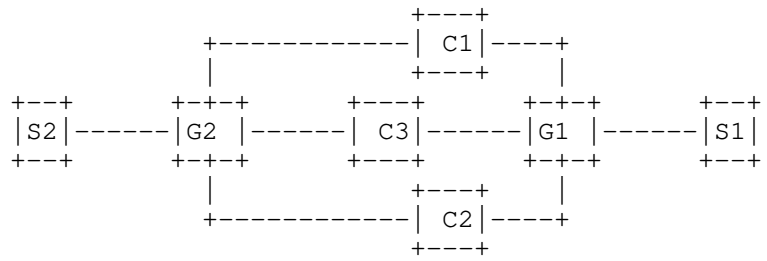

Figure 9: Multiple DOTS Clients, Multiple DOTS Gateways, Multiple DOTS Servers

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 decided 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 needs 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 place 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 and Nik Teague 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., 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>.

[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-04 (work in progress), April 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-22 (work in progress), September 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.

[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: 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.

Table of Contents

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]):

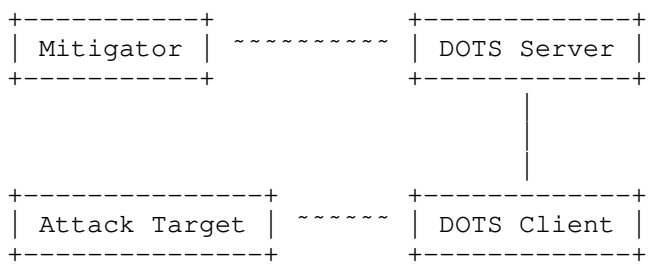

Figure 1: Basic DOTS Architecture

[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:

- 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.
- DHCP refers to both DHCPv4 [RFC2131] and DHCPv6 [RFC3315].
- DHCP client denotes a node that initiates requests to obtain configuration parameters from one or more DHCP servers.
- DHCP server refers to a node that responds to requests from DHCP clients.
- DOTS client: A DOTS-aware software module responsible for requesting attack response coordination with other DOTS-aware elements.
- 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.
- 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:

- 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.

- 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.

- 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).

- 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.

- 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.

- 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:

- "_dots._signal._udp.local."
- "_dots._signal._tcp.local."
- "_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.

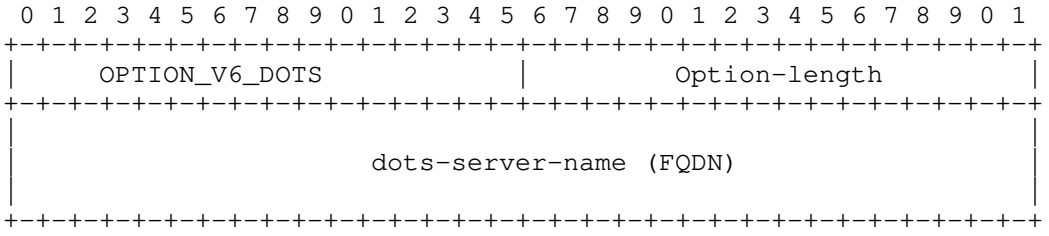

Figure 2: DHCPv6 DOTS Reference Identifier option

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

- Option-code: OPTION_V6_DOTS_RI (TBA1, see Section 12.1)
- Option-length: Length of the dots-server-name field in octets.
- 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.".

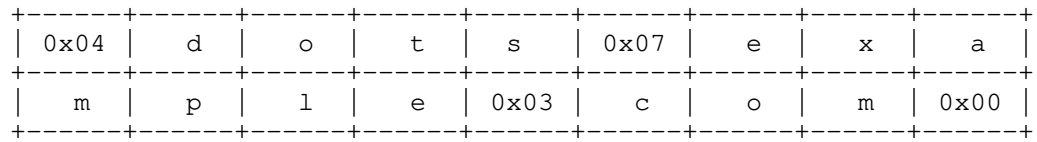

Figure 3: An example of the dots-server-name encoding

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.

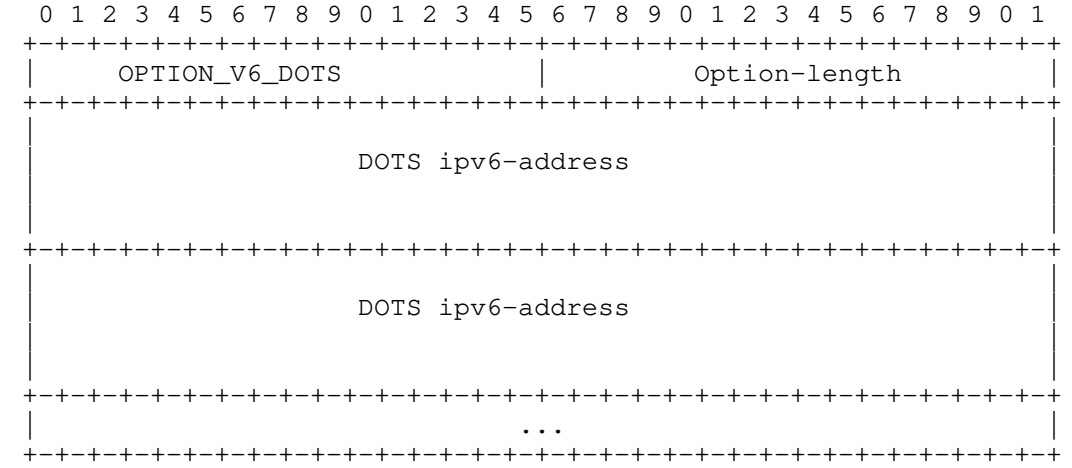

Figure 4: DHCPv6 DOTS Address option

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

- Option-code: OPTION_V6_DOTS_ADDRESS (TBA2, see Section 12.1)
- Option-length: Length of the 'DOTS ipv6-address(es)' field in octets. MUST be a multiple of 16.
- 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.

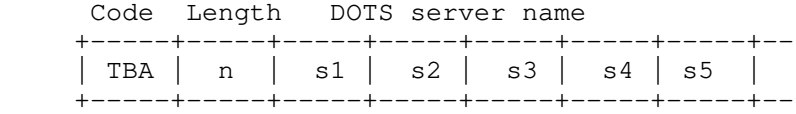

The values s1, s2, s3, etc. represent the domain name labels in the domain name encoding.

Figure 5: DHCPv4 DOTS Reference Identifier option

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

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

- 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 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.

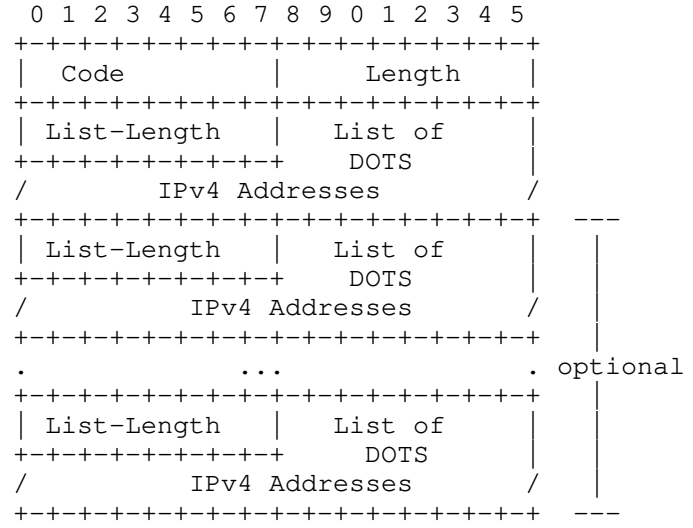

Figure 6: DHCPv4 DOTS Address option

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

- Code: OPTION_V4_DOTS_ADDRESS (TBA4, see Section 12.2);
- Length: Length of all included data in octets. The minimum length is 5.
- List-Length: Length of the "List of DOTS IPv4 Addresses" field in octets; MUST be a multiple of 4.
- 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.
- 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.

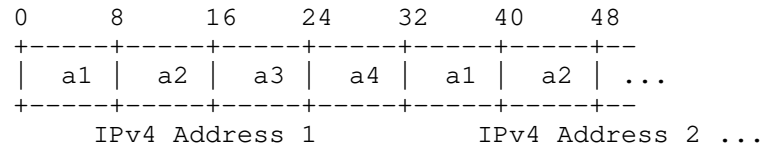

This format assumes that an IPv4 address is encoded as a1.a2.a3.a4.

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_V6_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

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

- Contact Information: <one of the authors>

12.3.2. signal.udp Application Protocol Tag Registration

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

12.3.3. signal.tcp Application Protocol Tag Registration

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

12.3.4. data.tcp Application Protocol Tag Registration

- Application Protocol Tag: data.tcp
- Intended Usage: See Section 6
- Security Considerations: See Section 11
- 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: praspati@cisco.com

DOTS
Internet-Draft
Intended status: Experimental
Expires: April 18, 2019

Y. Hayashi
NTT
K. Nishizuka
NTT Communications

October 15, 2018

DDoS mitigation offload usecase and YANG module expansion in signal channel
draft-h-dots-mitigation-offload-expansion-00

Abstract

This document describes a DDoS Mitigation offload usecase and an expansion of the YANG module in the DOTS signal channel for mitigating DDoS attack traffic correctly with general routers or switches. The proposed usecase and YANG module enhance DOTS capability to send attacker information and enable service providers to mitigate DDoS attack traffic by using general routers or switches in their intra-domain NW.

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 18, 2019.

Table of Contents

1. Introduction

Volume based distributed denial-of-service (DDoS) attacks such as DNS amplification attacks are threats for internet service providers because of their impact on network services. When such attacks occur, service providers have to mitigate them immediately to protect or recover their service. 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 concerned with DDoS attack detection, mitigation and so on collaborate.

On the other hand, the number of DDoS Mitigation Systems (DMS) that can be deployed in a service providers network is limited due to equipment cost. Thus, DMS's utilization rate can reach maximum capacity soon when the volume of DDoS attacks is enormous. When the rate reaches maximum capacity, the network needs to offload mitigation action from the DMS to cost-effective network devices such as switches and routers.

DDoS Open Threat Signaling (DOTS) is a protocol to standardize real-time signaling, threat-handling requests, and data between the multi-vendor elements [I-D.ietf-dots-use-cases]. This document describes an automated DDoS Mitigation offload usecase inherited from a DOTS usecase [I-D.ietf-dots-use-cases], which enables cost-effective DDoS Mitigation in an intra-domain network. Furthermore, this document describes an expansion of the YANG module in the DOTS signal channel

[I-D.ietf-dots-signal-channel], which enables a service provider's network to mitigate attack traffic correctly in the usecase.

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:

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 an inacceptable level.

DDoS attackers: Devices that carry out DDoS attacks.

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

Top Talker: A top N list of attackers who attack the same target. The list is ordered in terms of a two-tuple bandwidth such as bps or pps.

3. DDoS Mitigation Offload Usecase

The purpose of this usecase is to protect intra-domain network from volume-based DDoS attacks automatically, cost-effectively, and vendor-independently. The usecase is inherited from the DDoS Orchestration usecase in [I-D.ietf-dots-use-cases] and works on an intra-domain network.

Figure 1 and Figure 2 show a component diagram and C-plane sequence diagram of the usecase, respectively.

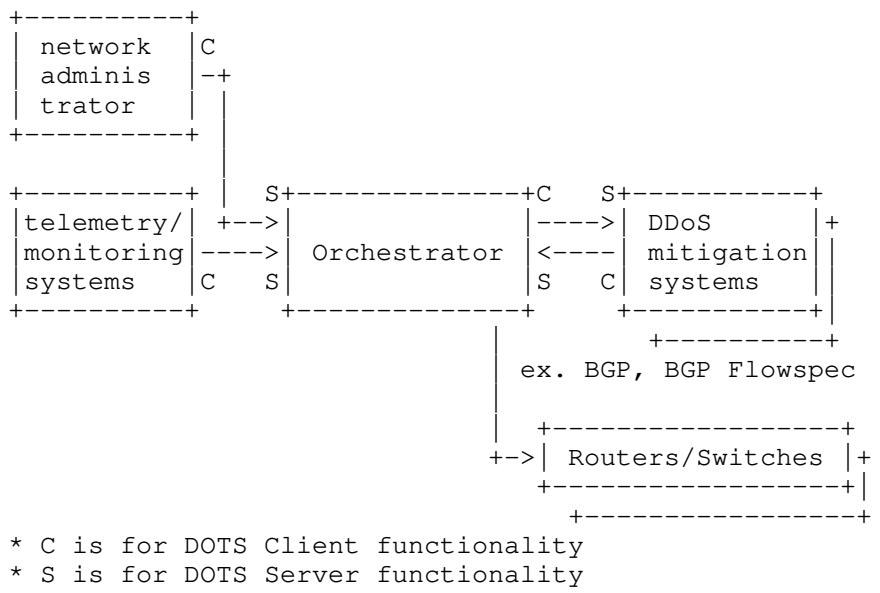

Figure 1: Component diagram of DDoS Mitigation offload usecase

This 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 DDoS mitigation systems have a DOTS client function to send mitigation requests to the orchestrator. Second, the orchestrator sends a request to routers or switches to block attack traffic.

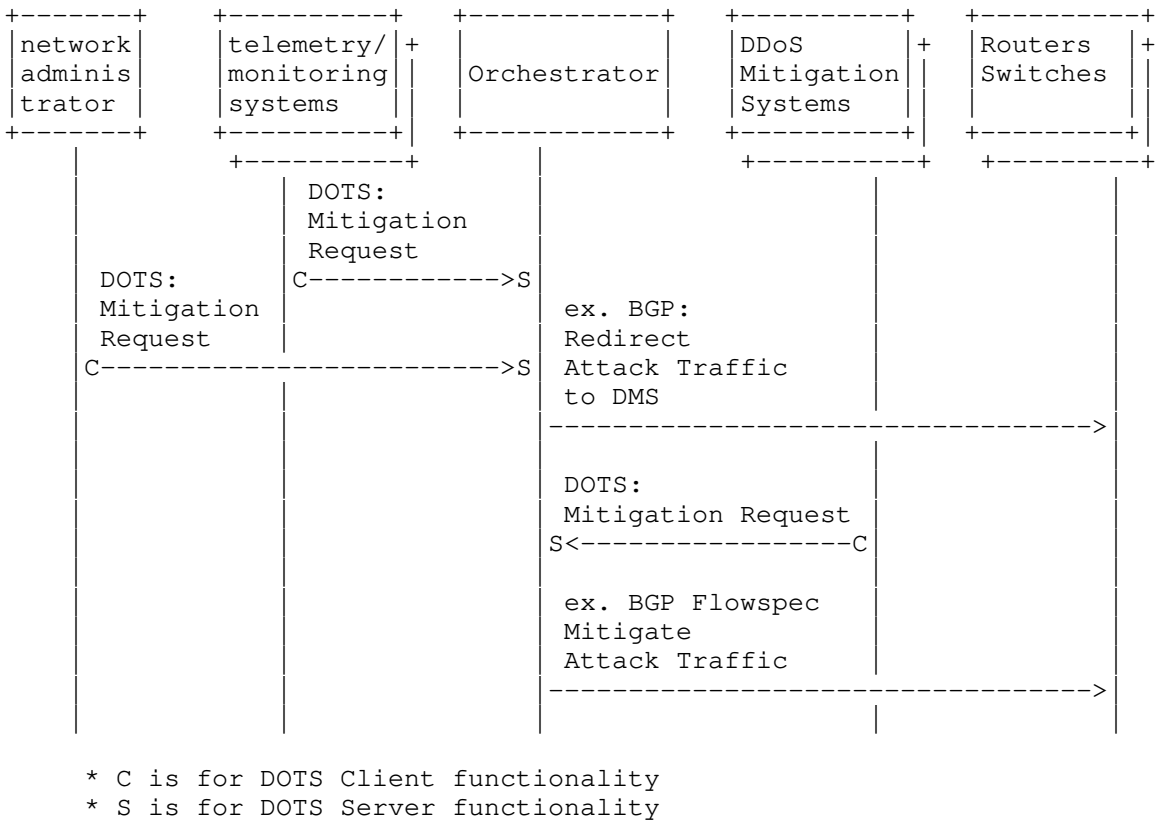

Figure 2: C-plane Sequence diagram of DDoS Mitigation offload usecase

In this usecase, when the telemetry/monitoring system detects a volume-based DDoS attack in the network, it sends a DOTS mitigation request to the orchestrator with target information such as target-prefix. Then, the network administrator confirms the request and sends a DOTS mitigation request to the orchestrator with the target information.

After that, the orchestrator requests the routers or switches to redirect attack traffic to the DMS by a configuration protocol such as a routing protocol like BGP [RFC4271] on the basis of the target information. Then the DMS analyzes attack traffic in detail and detects not only target but also attacker information, such as top-talker, and mitigates the attack traffic on the basis of the detected information.

When the volume-based attack becomes intense, DMS's utilization rate can reach maximum capacity. Then the DMS sends a DOTS mitigation request to the orchestrator as an offload request with the detection information. After that, the orchestrator requests the routers or switches to block attack traffic to the DMS by dissemination of flow specification rules protocols such as BGP flowspec [RFC5575] on the basis of the detected information.

4. Expansion of DOTS Signal Channel

It is desirable that the routers or switches mitigate attack traffic correctly after the DMS sends a DOTS Mitigation Request as an offload request in the usecase described in Section 3. For mitigating attack traffic correctly, this document proposes expanding DOTS signal channel [I-D.ietf-dots-signal-channel] so that it can send not only target information but also representative attacker information such as top talker. Note that it is difficult to send all attacker information because there is an enormous number of attackers when a volume-based DDoS attack occurs.

This section describes expansion of the YANG module [RFC7950] and mapping parameters to CBOR [RFC7049] of the DOTS Signal Channel.

4.1. Expansion of YANG Module of DOTS Signal Channel

Figure 3 shows an expanded YANG Module of the DOTS Signal Channel. Note that the "augment" statement allows a module to insert additional nodes into existing data models. The module defines a new grouping "attacker" and adds the grouping to an existing Signal Channel module by using an "augment" statement.

```
module ietf-dots-signal-channel-mitigation-offload-expansion {
     yang-version 1.1;

     namespace "urn:ietf:params:xml:ns:yang:
                ietf-dots-signal-channel:mitigation-offload-expansion";

     import ietf-dots-signal-channel {
       prefix signal;
     }

     import ietf-inet-types {
       prefix inet;
     }

     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:  Yuhei Hayashi
                 <mailto:hayashi.yuhei@lab.ntt.co.jp>
     description
       "This module contains the YANG definition for expanding signaling
        messages exchanged between a DOTS client and a DOTS server.

        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 2018-07-30 {
       description
         "Initial revision.";
       reference
         "RFC XXXX: ietf-dots-signal-channel";
     }

    /*
     * Groupings
     */
     grouping attacker {
       description
         "Specifies the attackers of the mitigation request.";
       leaf-list attacker-top-talker-prefix {
         type inet:ip-prefix;
         description
           "IPv4/IPv6 prefix identifying the top-talker in attackers.";
       }
     }

    /*
     * Main Container for DOTS Signal Channel Expansion
     */
     augment "/signal:dots-signal/signal:scope/"{
       uses attacker;
     }
```

```
}
```

Figure 3: Expansion of YANG Module of DOTS Signal Channel

4.2. Expansion of Mapping Parameters to CBOR

Figure 4 shows expansion of Mapping Parameters to CBOR [RFC7049] related to Figure 3.

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
...	...	...	...	...
attacker-top-talker-prefix	leaf-list inet: ip-prefix	XX	4 array 3 text string	Array String

Figure 4: Expansion of Mapping Parameters to CBOR

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>.

[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>.

[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., Moskowitz, R., and R. K, "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", draft-ietf-dots-requirements-15 (work in progress), August 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.

Authors' Addresses

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

Email: hayashi.yuhei@lab.ntt.co.jp, yuuhei.hayashi@gmail.com

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

Email: kaname@nttv6.jp

DOTS
Internet-Draft
Intended status: Standards Track
Expires: April 19, 2019

T. Reddy
J. Harsha
McAfee
M. Boucadair
Orange
J. Shallow
NCC Group

October 16, 2018

Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Call Home draft-reddy-dots-home-network-00

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 April 19, 2019.

Table of Contents

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 are 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 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 attack, 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 do 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. 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.

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 mitigation requests whenever the ISP detects there is an attack from a compromised device in the DOTS server's domain.

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 sample Call Home flow exchange:

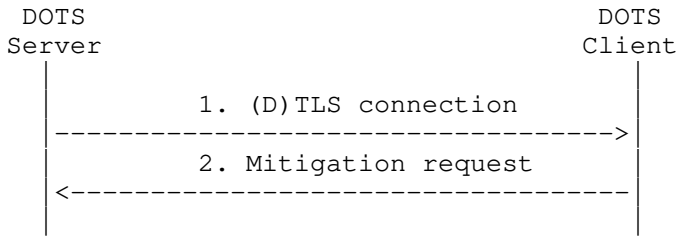

Figure 1: DOTS Signal Channel Call Home Sequence Diagram

This diagram makes the following points:

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

This is an optional attribute.

The 'source-prefix' and 'target-prefix' parameters are mandatory attributes when the attack traffic information is signaled by the DOTS client. 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.

The DOTS server uses the attack traffic information to find the pre-NAT 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. 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.

TBD:

a) Do we also want to convey Attack Name/type or ID (the home router may not be capable of detecting new emerging/sophisticated attacks) ?

b) Is DOTS data channel Call Home service required (if required, can RESTCONF Call Home defined in RFC8071 be used) ?

3.2.2. DOTS Signal Call Home YANG Module

3.2.2.1. Mitigation Request 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 structure:

```
module: ietf-dots-signal-call-home
  augment /ietf-signal:dots-signal:
    +--rw source-prefix*       inet:ip-prefix
    +--rw source-port-range* [lower-port upper-port]
       +--rw lower-port    inet:port-number
       +--rw upper-port    inet:port-number
    +--rw source-icmp-type-range* [lower-type upper-type]
       +--rw lower-type    uint8
       +--rw upper-type    uint8
```

3.2.2.2. Call Home Mitigation Request 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 XXXX: 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:jon.shallow@nccgroup.com>";

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

         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";
      }

      augment "/ietf-signal:dots-signal" {
        when "message-type='mitigation-scope'";
        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 number.";
        leaf lower-type {
          type uint8;
          mandatory true;
          description
            "Lower ICMP type number of the ICMP type range.";
        }
        leaf upper-type {
          type uint8;
          must ". >= ../lower-type" {
             error-message
               "The upper ICMP type number must be greater than
               or equal to lower ICMP type number.";
          }
          description
            "Upper type number of the ICMP type range.";
        }
      }
    }
  }
```

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 and source-port-range are comprehension-optional parameters.

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

Table 4: CBOR Mappings Used in DOTS Signal Channel Messages

4.3. DOTS Signal Channel YANG Module

This document requests IANA to register the following URIs 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 modules 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.

6. Acknowledgements

TBC.

7. References

7.1. 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>.

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

[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>.

7.2. Informative References

[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-15 (work in progress), August 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.

[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>.

[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>.

[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.

Table of Contents

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.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:

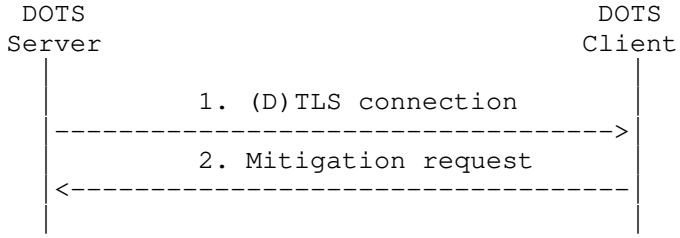

Figure 1: DOTS Signal Channel Call Home Sequence Diagram

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.

- 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 3 text string	Array String
source-port-range	list	0x8001 (TBD2)	4 array	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