

Internet Engineering Task Force (IETF)
Request for Comments: 9066
Category: Standards Track
ISSN: 2070-1721

T. Reddy.K
Akamai
M. Boucadair, Ed.
Orange
J. Shallow
December 2021

Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Call Home

Abstract

This document specifies the Denial-of-Service Open Threat Signaling (DOTS) signal channel Call Home, which enables a Call Home DOTS server to initiate a secure connection to a Call Home DOTS client and to receive attack traffic information from the Call Home DOTS client. The Call Home DOTS server in turn uses the attack traffic information to identify compromised devices launching outgoing DDoS attacks and take appropriate mitigation action(s).

The DOTS signal channel Call Home is not specific to home networks; the solution targets any deployment in which it is required to block DDoS attack traffic closer to the source(s) of a DDoS attack.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9066>.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
2. Terminology
3. Applicability Scope
4. Coexistence of a Base DOTS Signal Channel and DOTS Call Home
5. DOTS Signal Channel Call Home
 - 5.1. Procedure

- 5.2. DOTS Signal Channel Variations
 - 5.2.1. Heartbeat Mechanism
 - 5.2.2. Redirected Signaling
- 5.3. DOTS Signal Channel Extension
 - 5.3.1. Mitigation Request
 - 5.3.2. Address Sharing Considerations
- 6. DOTS Signal Call Home YANG Module
 - 6.1. Tree Structure
 - 6.2. YANG/JSON Mapping Parameters to CBOR
 - 6.3. YANG Module
- 7. IANA Considerations
 - 7.1. DOTS Signal Channel CBOR Mappings Registry
 - 7.2. New DOTS Conflict Cause
 - 7.3. DOTS Signal Call Home YANG Module
- 8. Security Considerations
- 9. Privacy Considerations
- 10. References
 - 10.1. Normative References
 - 10.2. Informative References
- Appendix A. Some Home Network Issues
- Appendix B. Disambiguating Base DOTS Signal vs. DOTS Call Home
- Acknowledgements
- Contributors
- Authors' Addresses

1. Introduction

The Distributed Denial-of-Service Open Threat Signaling (DOTS) signal channel protocol [RFC9132] 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 [RFC4732]. 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 [RFC8903].

However, [RFC9132] only covers how to mitigate when being attacked (i.e., protecting a network from inbound DDoS attacks). It does not cover how to control the attacks close to their source(s) that are misusing network resources (i.e., outbound DDoS attacks). In particular, the DOTS signal protocol does not discuss cooperative DDoS mitigation between the network hosting an attack source and the Internet Service Provider (ISP) to suppress the outbound DDoS attack traffic originating from that network. As a reminder, the base basic DOTS architecture is depicted in Figure 1 (Section 2 of [RFC8811]).

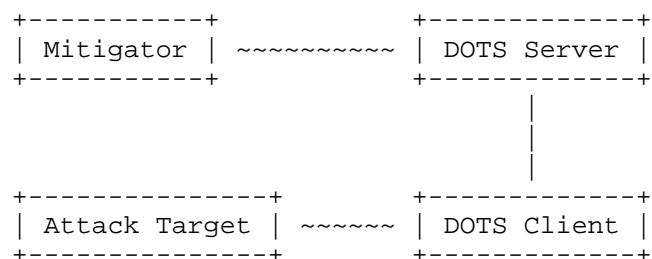


Figure 1: Basic DOTS Architecture

Appendix A details why the rise of Internet of Things (IoT) compounds the possibility of these being used as malicious actors that need to be controlled. Similar issues can be encountered in enterprise networks, data centers, etc. The ISP offering a DDoS mitigation service can detect outgoing DDoS attack traffic originating from its subscribers, or the ISP may receive filtering rules (e.g., using BGP Flowspec [RFC8955] [RFC8956]) from a transit provider to filter, block, or rate-limit DDoS attack traffic originating from the ISP's subscribers to a downstream target. Nevertheless, the DOTS signal

channel does not provide means for the ISP to request blocking such attacks close to the sources without altering legitimate traffic. This document fills that void by specifying an extension to the DOTS signal channel: DOTS signal channel Call Home.

Note: Another design approach would be to extend the DOTS signal channel with a new attribute to explicitly indicate whether a mitigation request concerns an outbound DDoS attack. In such an approach, it is assumed that a DOTS server is deployed within the domain that is hosting the attack source(s), while a DOTS client is enabled within an upstream network (e.g., access network). However, initiating a DOTS signal channel from an upstream network to a source network is complicated because of the presence of translators and firewalls. Moreover, the use of the same signal channel to handle both inbound and outbound attacks complicates both the heartbeat and redirection mechanisms that are executed as a function of the attack direction (see Sections 5.2.1 and 5.2.2). Also, the DOTS server will be subject to fingerprinting (e.g., using scanning tools) and DoS attacks (e.g., by having the DOTS server perform computationally expensive operations). Various management and deployment considerations that motivate the Call Home functionality are listed in Section 1.1 of [RFC8071].

"DOTS signal channel Call Home" (or "DOTS Call Home" for short) refers to a DOTS signal channel established at the initiative of a DOTS server thanks to a role reversal at the (D)TLS layer (Section 5.1). That is, the DOTS server initiates a secure connection to a DOTS client and uses that connection to receive the attack traffic information (e.g., attack sources) from the DOTS client.

A high-level DOTS Call Home functional architecture is shown in Figure 2. Attack source(s) are within the DOTS server domain.

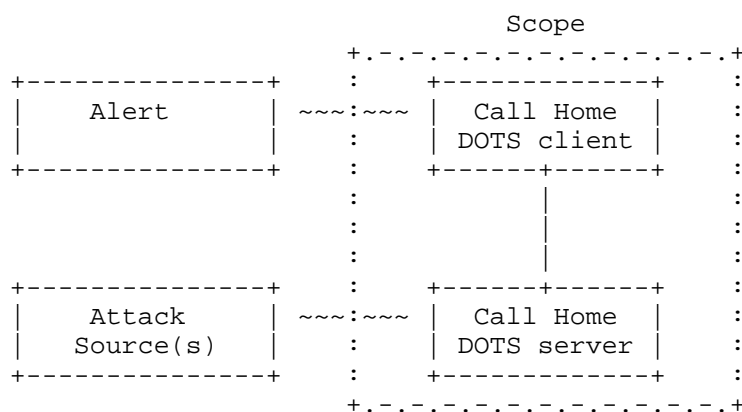


Figure 2: Basic DOTS Signal Channel Call Home Functional Architecture

DOTS agents involved in the DOTS Call Home otherwise adhere to the DOTS roles as defined in [RFC8612]. For clarity, this document uses "Call Home DOTS client" (or "Call Home DOTS server") to refer to a DOTS client (or DOTS server) deployed in a Call Home scenario (Figure 2). Call Home DOTS agents may (or may not) be co-located with DOTS agents that are compliant with [RFC9132] (see Section 4 for more details).

A Call Home DOTS client relies upon a variety of triggers to make use of the Call Home function (e.g., scrubbing the traffic from the attack source or receiving an alert from an attack target, a peer DDoS Mitigation System (DMS), or a transit provider). The definition of these triggers is deployment specific. It is therefore out of the scope of this document to elaborate on how these triggers are made available to a Call Home DOTS client.

In a typical deployment scenario, the Call Home DOTS server is enabled on a Customer Premises Equipment (CPE), which is aligned with recent trends to enrich the CPE with advanced security features. For example, the DOTS Call Home service can be part of services supported by an ISP-managed CPE or a managed security service subscribed to by the user. Unlike classic DOTS deployments [RFC8903], a Call Home DOTS server maintains a single DOTS signal channel session for each DOTS-capable upstream provisioning domain [DOTS-MULTIHOMING].

For instance, the Call Home DOTS server in the home network initiates the signal channel Call Home in "idle" time; subsequently, the Call Home 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 (i.e., from within the home network).

The Call Home DOTS server uses the DDoS attack traffic information to identify the compromised device in its domain that is responsible for launching the DDoS attack, optionally notifies a network administrator, and takes appropriate mitigation action(s). For example, a mitigation action can be to quarantine the compromised device or block its traffic to the attack target(s) until the mitigation request is withdrawn.

This document assumes that Call Home DOTS servers are provisioned with a way to know how to reach the upstream Call Home DOTS client(s), which could occur by a variety of means (e.g., [RFC8973]). The specification of such means are out of scope of this document.

More information about the applicability scope of the DOTS signal channel Call Home is provided in Section 3.

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 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 Section 1.2 of [RFC8612].

"DDoS Mitigation System (DMS)" refers to a system that performs DDoS mitigation.

"Base DOTS signal channel" refers to [RFC9132].

The meaning of the symbols in YANG tree diagrams are defined in [RFC8340] and [RFC8791].

(D)TLS is used for statements that apply to both Transport Layer Security (TLS) [RFC8446] and Datagram Transport Layer Security (DTLS) [RFC6347] [DTLS13]. Specific terms are used for any statement that applies to either protocol alone.

3. Applicability Scope

The problems discussed in Section 1 may be encountered in many deployments (e.g., home networks, enterprise networks, transit networks, data centers). 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. That is, attacks that are issued, e.g., from within an enterprise network or a data center will thus be blocked before exiting these networks.

An instantiation of the Call Home functional architecture is depicted in Figure 3.

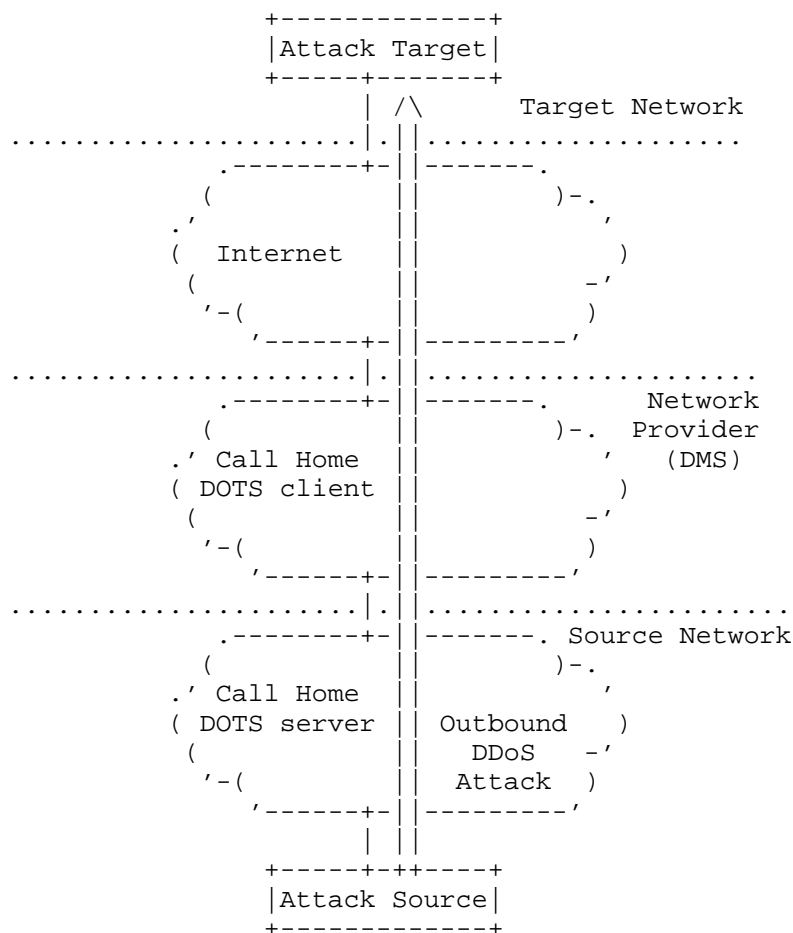


Figure 3: DOTS Signal Channel Call Home Reference Architecture

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

Call Home DOTS agent relationships are similar to those discussed in Section 2.3 of [RFC8811]. For example, multiple Call Home DOTS servers of the same domain can be associated with the same Call Home DOTS client. A Call Home DOTS client may decide to contact these Call Home DOTS servers sequentially, fork a mitigation request to all of them, or select one Call Home DOTS server to place a mitigation request. Such a decision is implementation specific.

For some mitigations, feedback may be required from an administrator to confirm a filtering action. The means to seek an administrator's consent are deployment specific. Indeed, a variety of implementation options can be considered for any given Call Home DOTS deployment, such as push notifications using a dedicated application, Syslog, etc. It is out of the scope of this document to make recommendations about how such interactions are implemented (see Figure 2).

The Call Home DOTS server can be enabled on a border router or a dedicated appliance. For the particular case of home networks, the Call Home DOTS server functionality can be enabled on a managed CPE or bundled with a CPE management application that is provided by an ISP to its subscribers. These managed services are likely to be designed to hide the complexity of managing (including configuring) the CPE. For example, managed CPEs support the means to notify the user when a new device is detected in order to seek confirmation as

to whether or not access should be granted to the device. These means can be upgraded to interface with the Call Home DOTS server. Customized settings can be configured by users to control the notifications (e.g., triggers, type) and default actions.

4. Coexistence of a Base DOTS Signal Channel and DOTS Call Home

The DOTS signal channel Call Home does not require or preclude the activation of the base DOTS signal channel [RFC9132]. Some sample deployment schemes are discussed in this section for illustration purposes.

The network that hosts an attack source may also be subject to inbound DDoS attacks. In that case, both the base DOTS signal channel and DOTS signal channel Call Home may be enabled as shown in Figure 4 (same DMS provider) or Figure 5 (distinct DMS providers).

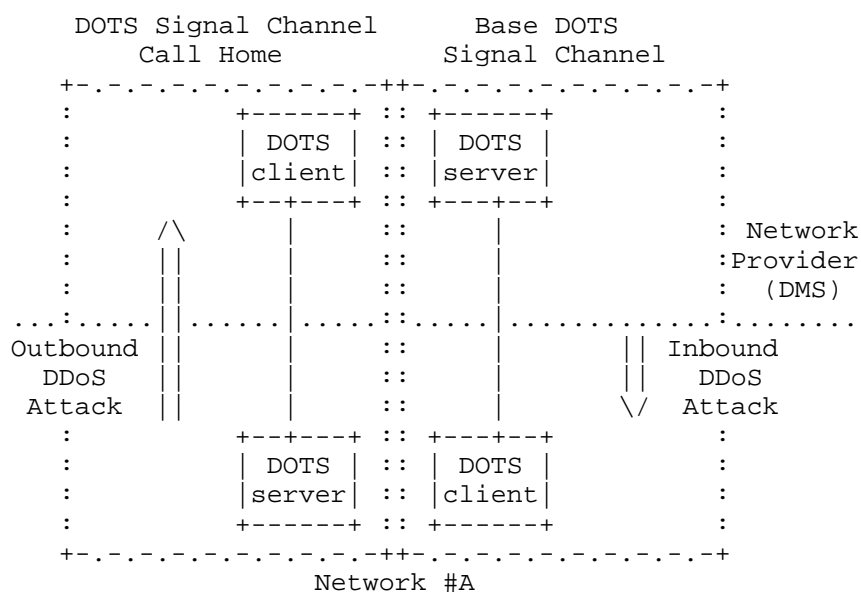
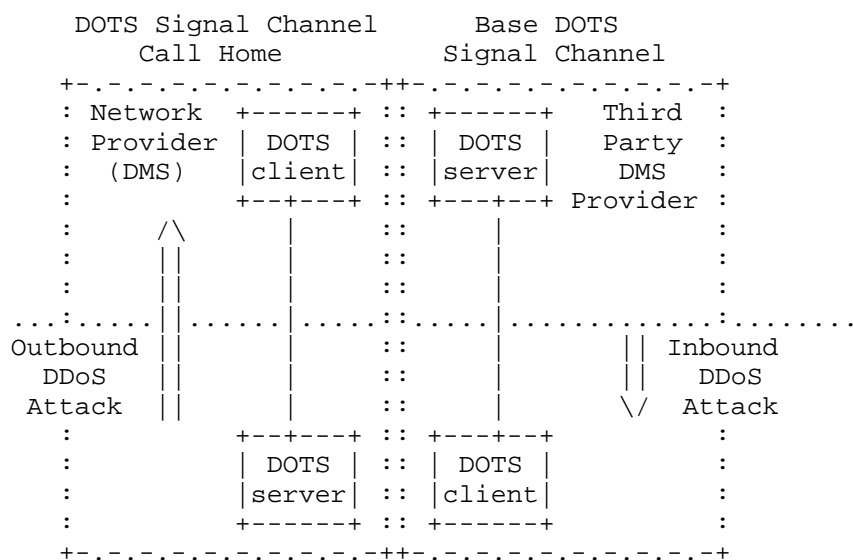


Figure 4: Activation of a Base DOTS Signal Channel and Call Home
(Same DMS Provider)

Note that a DMS provider may not be on the default forwarding path of inbound DDoS attack traffic targeting a network (e.g., Network #B in Figure 5). Nevertheless, the DOTS signal channel Call Home requires the DMS provider to be on the default forwarding path of the outbound traffic from a given network.



Network #B

Figure 5: Activation of a Base DOTS Signal Channel and Call Home (Distinct DMS Providers)

Figures 6 and 7 depict examples where the same node embeds both base DOTS and Call Home DOTS agents. For example, a DOTS server and a Call Home DOTS client may be enabled on the same device within the infrastructure of a DMS provider (e.g., Node #i in Figure 6), or a DOTS client and a Call Home DOTS server may be enabled on the same device within a source network (e.g., Node #j with Network #D shown in Figure 7).

Whether the same or distinct nodes are used to host base DOTS and Call Home DOTS agents is specific to each domain.

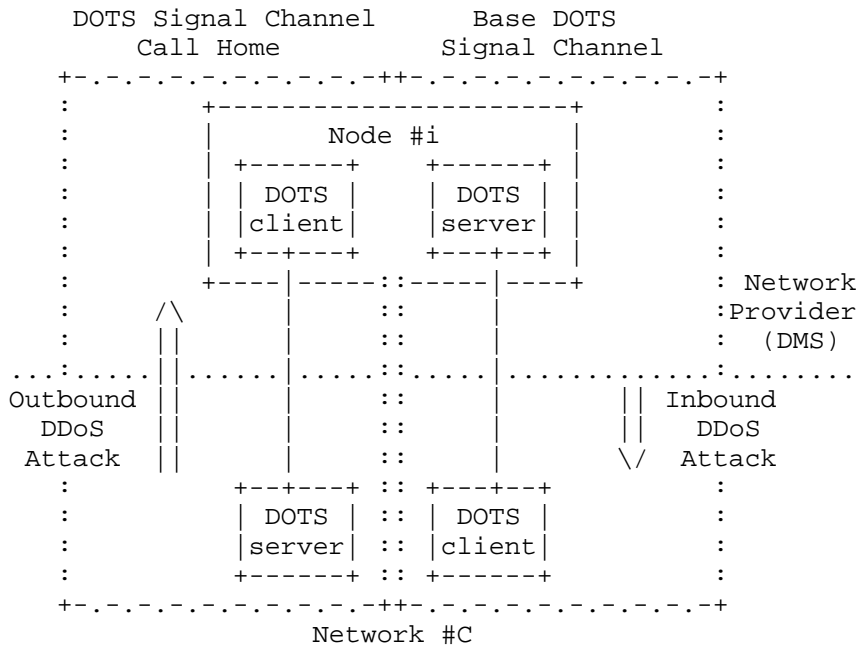
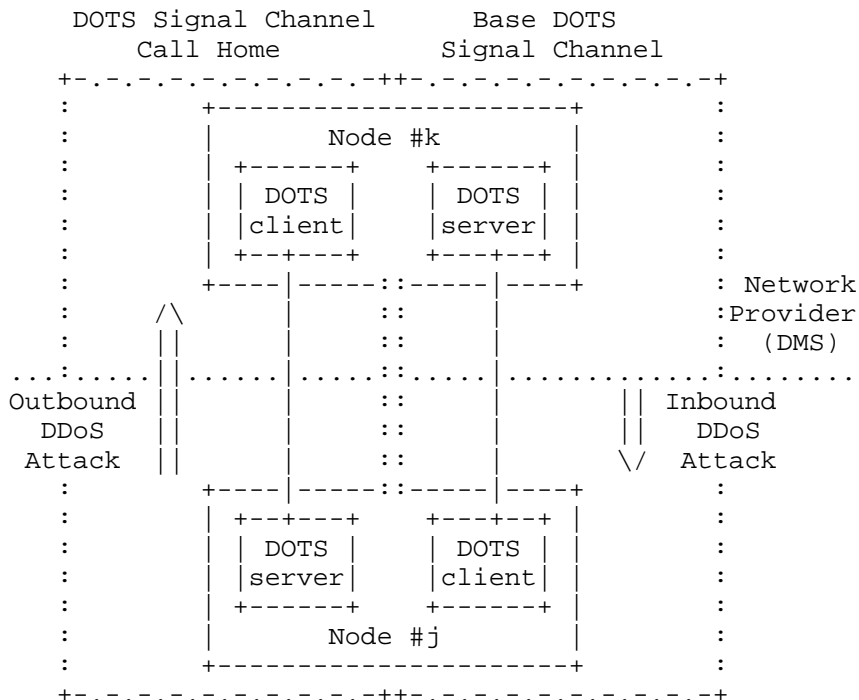


Figure 6: The Same Node Embedding a Call Home DOTS Client and a DOTS Server at the Network Provider's Side



Network #D

Figure 7: The Same Node Embedding both a DOTS Client and a Call Home DOTS Server

Appendix B elaborates on the considerations to unambiguously distinguish DOTS messages that belong to each of these channels.

5. DOTS Signal Channel Call Home

5.1. Procedure

The DOTS signal channel Call Home preserves all but one of the DOTS client/server roles in the DOTS protocol stack, as compared to the client-initiated DOTS signal channel protocol [RFC9132]. The role reversal that occurs is at the (D)TLS layer; that is, (1) the Call Home DOTS server acts as a DTLS client, and the Call Home DOTS client acts as a DTLS server; or (2) the Call Home DOTS server acts as a TLS client initiating the underlying TCP connection, and the Call Home DOTS client acts as a TLS server. The Call Home DOTS server initiates a (D)TLS handshake to the Call Home DOTS client.

For example, a home network element (e.g., home router) co-located with a Call Home DOTS server is the (D)TLS client. That is, the Call Home DOTS server assumes the role of the (D)TLS client, but the network element's role as a DOTS server remains the same.

Existing certificate chains and mutual authentication mechanisms between the DOTS agents are unaffected by the Call Home function. From a deployment standpoint, and given the scale of Call Home DOTS servers that may be involved, enabling means for automating the provisioning of credentials on Call Home DOTS servers to authenticate to the Call Home DOTS client is encouraged. It is out of the scope of this document to elaborate on these means.

Figure 8 illustrates a sample DOTS Call Home flow exchange:

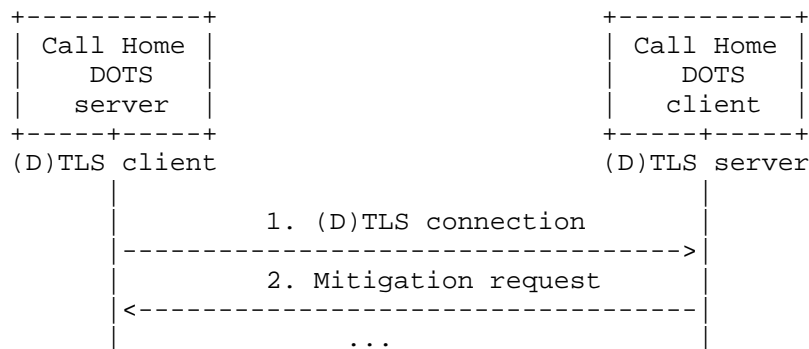


Figure 8: DOTS Signal Channel Call Home Sequence Diagram

The DOTS signal channel Call Home procedure is as follows:

1. If UDP transport is used, the Call Home DOTS server begins by initiating a DTLS connection to the Call Home DOTS client.

If TCP is used, the Call Home DOTS server begins by initiating a TCP connection to the Call Home DOTS client. Once connected, the Call Home DOTS server continues to initiate a TLS connection to the Call Home DOTS client.

Peer DOTS agents may have mutual agreement to use a specific port number, such as by explicit configuration or dynamic discovery [RFC8973]. The interaction between the base DOTS signal channel and the Call Home is discussed in Appendix B.

The Happy Eyeballs mechanism explained in Section 4.3 of [RFC9132] is used for initiating (D)TLS connections.

2. Using this (D)TLS connection, the Call Home DOTS client may request, withdraw, or retrieve the status of mitigation requests. The Call Home DOTS client supplies the source information by means of new attributes defined in Section 5.3.1.

The heartbeat mechanism used for the DOTS Call Home deviates from the one defined in Section 4.7 of [RFC9132]. Section 5.2.1 specifies the behavior to be followed by Call Home DOTS agents.

5.2. DOTS Signal Channel Variations

5.2.1. Heartbeat Mechanism

Once the (D)TLS session is established between the DOTS agents, the Call Home DOTS client contacts the Call Home DOTS server to retrieve the session configuration parameters (Section 4.5 of [RFC9132]). The Call Home DOTS server adjusts the "heartbeat-interval" to accommodate binding timers used by on-path NATs and firewalls. Heartbeats will then be exchanged by the DOTS agents following the instructions retrieved using the signal channel session configuration exchange.

It is the responsibility of Call Home DOTS servers to ensure that on-path translators/firewalls are maintaining a binding so that the same external IP address and/or port number is retained for the DOTS signal channel session. A Call Home DOTS client MAY trigger their heartbeat requests immediately after receiving heartbeat probes from its peer Call Home DOTS server.

When an outgoing attack that saturates the outgoing link from the Call Home DOTS server is detected and reported by a Call Home DOTS client, the latter MUST continue to use the DOTS signal channel even if no traffic is received from the Call Home DOTS server.

If the Call Home DOTS server receives traffic from the Call Home DOTS client, the Call Home DOTS server MUST continue to use the DOTS signal channel even if the threshold of allowed missing heartbeats ("missing-hb-allowed") is reached.

If the Call Home DOTS server does not receive any traffic from the peer Call Home DOTS client during the time span required to exhaust the maximum "missing-hb-allowed" threshold, the Call Home DOTS server concludes the session is disconnected. Then, the Call Home DOTS server MUST try to establish a new DOTS signal channel session, preferably by resuming the (D)TLS session.

5.2.2. Redirected Signaling

A Call Home DOTS server MUST NOT support the redirected signaling mechanism as specified in Section 4.6 of [RFC9132] (i.e., a 5.03 response that conveys an alternate DOTS server's Fully Qualified Domain Name (FQDN) or IP address(es)). A Call Home DOTS client MUST silently discard such a message as only a Call Home DOTS server can initiate a new (D)TLS connection.

If a Call Home DOTS client wants to redirect a Call Home DOTS server to another Call Home DOTS client, it MUST send a Non-confirmable PUT request to the predefined resource ".well-known/dots/redirect" with the following attributes in the body of the PUT request:

alt-ch-client: The FQDN of an alternate Call Home DOTS client. It

is also presented as a reference identifier for authentication purposes.

This is a mandatory attribute for DOTS signal Call Home. It MUST NOT be used for base DOTS signal channel operations.

alt-ch-client-record: List of IP addresses for the alternate Call Home DOTS client. If no "alt-ch-client-record" is provided, the Call Home DOTS server passes the "alt-ch-client" name to a name resolution library to retrieve one or more IP addresses of the alternate Call Home DOTS client.

This is an optional attribute for DOTS signal Call Home. It MUST NOT be used for base DOTS signal channel operations.

ttl: The Time To Live (TTL) of the alternate Call Home DOTS client. That is, the time interval in which the alternate Call Home DOTS client may be cached for use by a Call Home DOTS server.

This is an optional attribute for DOTS signal Call Home. It MUST NOT be used for base DOTS signal channel operations.

On receipt of this PUT request, the Call Home DOTS server responds with a 2.01 (Created), closes this connection, and establishes a connection with the new Call Home DOTS client. The processing of the TTL is defined in Section 4.6 of [RFC9132]. If the Call Home DOTS server cannot service the PUT request, the response is rejected with a 4.00 (Bad Request).

Figure 9 shows a PUT request example to convey the alternate Call Home DOTS client "alt-call-home-client.example" together with its IP addresses 2001:db8:6401::1 and 2001:db8:6401::2. The validity of this alternate Call Home DOTS client is 10 minutes.

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "redirect"
Uri-Path: "cuid=dz6pHjaADkaFTbjr0JGBpw"
Uri-Path: "mid=123"
Content-Format: "application/dots+cbor"

{
  "ietf-dots-signal-channel:redirected-signal": {
    "ietf-dots-call-home:alt-ch-client":
      "alt-call-home-client.example",
    "ietf-dots-call-home:alt-ch-client-record": [
      "2001:db8:6401::1",
      "2001:db8:6401::2"
    ],
    "ietf-dots-call-home:ttl": 600
  }
}
```

Figure 9: Example of a PUT Request for Redirected Signaling

Figure 9 uses the JSON encoding of YANG-modeled data for the CoAP message body. The same encoding is used in Figure 10 (Section 5.3.1).

5.3. DOTS Signal Channel Extension

5.3.1. Mitigation Request

This specification extends the mitigation request defined in Section 4.4.1 of [RFC9132] to convey the attack source information (e.g., source prefixes, source port numbers). The DOTS client

conveys the following new parameters in the Concise Binary Object Representation (CBOR) body of the mitigation request:

source-prefix: A list of attacker IP prefixes used to attack the target. Prefixes are represented using Classless Inter-Domain Routing (CIDR) notation (BCP 122 [RFC4632]).

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

The prefix list MUST NOT include broadcast, loopback, or multicast addresses. These addresses are considered invalid values. Note that link-local addresses are allowed. The Call Home DOTS client MUST validate that attacker prefixes are within the scope of the Call Home DOTS server domain (e.g., prefixes assigned to the Call Home DOTS server domain or networks it services). This check is meant to avoid contacting Call Home DOTS servers that are not entitled to enforce actions on specific prefixes.

This is an optional attribute for the base DOTS signal channel operations.

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 any subrange of 0-65535 -- for example, 0-1023, 1024-65535, or 1024-49151.

This is an optional attribute for the base DOTS signal channel operations.

source-icmp-type-range: 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. Both ICMP [RFC0792] and ICMPv6 [RFC4443] types are supported. Whether ICMP or ICMPv6 types are to be used is determined by the address family of the "target-prefix".

This is an optional attribute for the base DOTS signal channel operations.

The "source-prefix" parameter is a mandatory attribute when the attack traffic information is signaled by a Call Home DOTS client (i.e., the Call Home scenario depicted in Figure 8). The "target-prefix" attribute MUST be included in the mitigation request signaling the attack information to a Call Home DOTS server. The "target-uri" or "target-fqdn" parameters can be included in a mitigation request for diagnostic purposes to notify the Call Home DOTS server domain administrator but SHOULD NOT be used to determine the target IP addresses. "alias-name" is unlikely to be conveyed in a Call Home mitigation request given that a target may be any IP resource and that there is no incentive for a Call Home DOTS server (embedded, for example, in a CPE) to maintain aliases.

In order to help attack source identification by a Call Home DOTS server, the Call Home DOTS client SHOULD include in its mitigation request additional information such as "source-port-range" or "source-icmp-type-range" to disambiguate nodes sharing the same "source-prefix". IPv6 addresses/prefixes are sufficient to uniquely

identify a network endpoint, without need for port numbers or ICMP type information. While this is also possible for IPv4, it is much less often the case than for IPv6. More address sharing implications on the setting of source information ("source-prefix", "source-port-range") are discussed in Section 5.3.2.

Only immediate mitigation requests (i.e., "trigger-mitigation" set to "true") are allowed; Call Home DOTS clients MUST NOT send requests with "trigger-mitigation" set to "false". Such requests MUST be discarded by the Call Home DOTS server with a 4.00 (Bad Request).

An example of a mitigation request sent by a Call Home DOTS client is shown in Figure 10.

```
Header: PUT (Code=0.03)
Uri-Path: ".well-known"
Uri-Path: "dots"
Uri-Path: "mitigate"
Uri-Path: "cuid=dz6pHjaADkaFTbjr0JGBpw"
Uri-Path: "mid=56"
Content-Format: "application/dots+cbor"

{
  "ietf-dots-signal-channel:mitigation-scope": {
    "scope": [
      {
        "target-prefix": [
          "2001:db8:c000::/128"
        ],
        "ietf-dots-call-home:source-prefix": [
          "2001:db8:123::1/128"
        ],
        "lifetime": 3600
      }
    ]
  }
}
```

Figure 10: An Example of a Mitigation Request Issued by a Call Home DOTS Client

The Call Home DOTS server MUST check that the "source-prefix" is within the scope of the Call Home DOTS server domain. Note that in a DOTS Call Home scenario, the Call Home DOTS server considers, by default, that any routable IP prefix enclosed in "target-prefix" is within the scope of the Call Home DOTS client. Invalid mitigation requests are handled as per Section 4.4.1 of [RFC9132].

Note: These validation checks do not apply when the source information is included as a hint in the context of the base DOTS signal channel.

Call Home 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 consent from the Call Home DOTS server domain administrator is required, the Call Home DOTS server replies with 2.01 (Created) and the "status" code set to 1 (attack-mitigation-in-progress). Then, the mechanisms defined in Section 4.4.2 of [RFC9132] are followed by the DOTS agents to update the mitigation status. In particular, if the attack traffic is blocked, the Call Home DOTS server informs the Call Home DOTS client that the attack is being mitigated (i.e., by setting the "status" code to 2 (attack-successfully-mitigated)).

If the attack traffic information is identified by the Call Home DOTS server or the Call Home DOTS server domain administrator as legitimate traffic, the mitigation request is rejected with a 4.09 (Conflict) (e.g., when no consent is required from an administrator) or a notification message with the "conflict-clause" (Section 4.4.1 of [RFC9132]) set to the following new value:

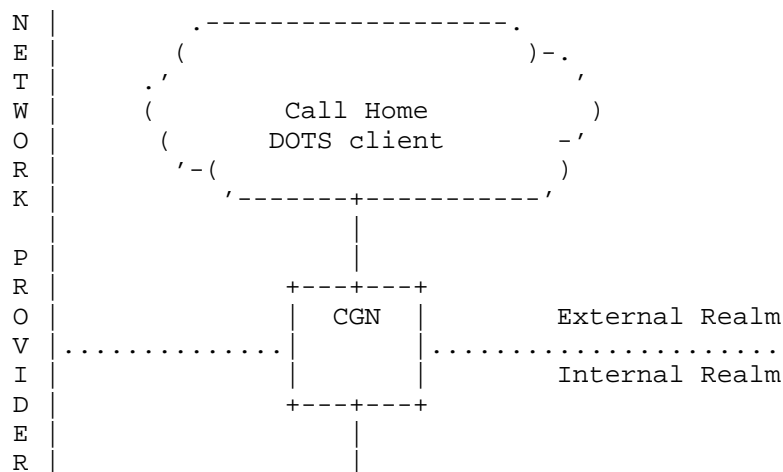
4: Mitigation request rejected. This code is returned by the DOTS server to indicate the attack traffic has been classified as legitimate traffic.

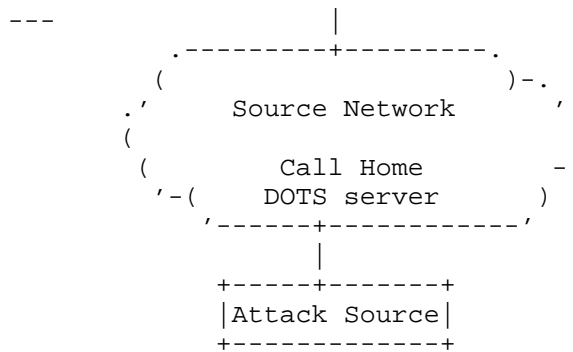
Once the request is validated by the Call Home DOTS server, appropriate actions are enforced to block the attack traffic within the source network. For example, if the Call Home DOTS server is embedded in a CPE, it can program the packet processor to punt all the traffic from the compromised device to the target to slow path. The CPE 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 CPE administrator about the compromised device. Note that the Call Home DOTS client is informed about the progress of the attack mitigation following the rules in Section 4.4.2 of [RFC9132].

The DOTS agents follow the same procedures specified in [RFC9132] for managing a mitigation request.

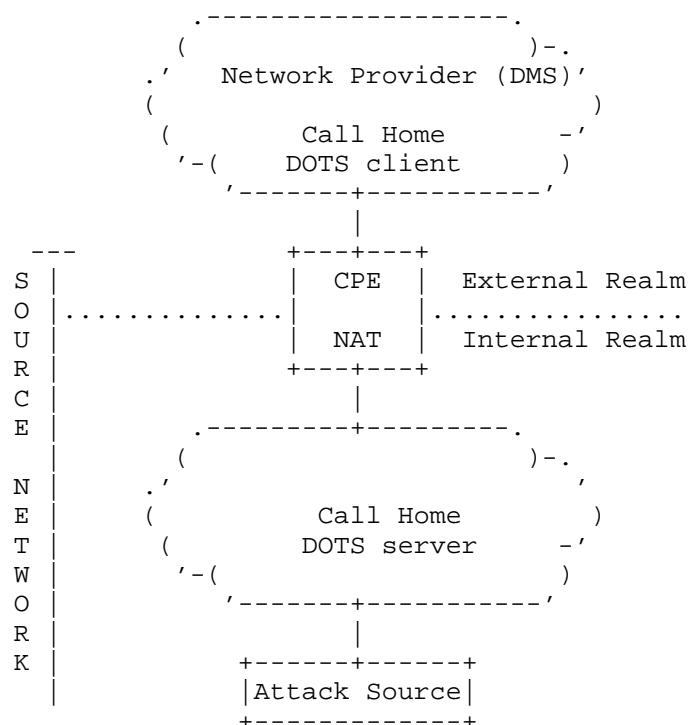
5.3.2. Address Sharing Considerations

Figure 11 depicts an example of a network provider that hosts a Call Home DOTS client and deploys a Carrier-Grade NAT (CGN) between the DOTS client domain and DOTS server domain. In such cases, communicating an external IP address in a mitigation request by a Call Home DOTS client is likely to be discarded by the Call Home DOTS server because the external IP address is not visible locally to the Call Home DOTS server (Figure 11). The Call Home DOTS server is only aware of the internal IP addresses/prefixes bound to its domain (i.e., those used in the internal realm shown in Figure 11). Thus, Call Home DOTS clients that are aware of the presence of on-path CGNs MUST NOT include the external IP address and/or port number identifying the suspect attack source (i.e., those used in the external realm shown in Figure 11) but MUST include the internal IP address and/or port number. To that aim, the Call Home DOTS client SHOULD rely on mechanisms, such as those described in [RFC8512] or [RFC8513], to retrieve the internal IP address and port number that are mapped to an external IP address and port number. For the particular case of NAT64 [RFC6146], if the target address is an IPv4 address, the IPv4-converted IPv6 address of this target address [RFC6052] SHOULD be used.





If a Mapping of Address and Port (MAP) Border Relay [RFC7597] or Lightweight Address Family Transition Router (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.



.....

Figure 13: Example of a Call Home DOTS Server and a NAT Embedded in a CPE

Note: Implementers must check that the mapping output provided by their YANG-to-CBOR encoding schemes is aligned with the content of Table 1.

Parameter Name	YANG Type	CBOR Key Value	CBOR Major Type & Information	JSON Type
ietf-dots-call-home:source-prefix	leaf-list inet:ip-prefix	32768	4 array 3 text string	Array String
ietf-dots-call-home:source-port-range	list	32769	4 array	Array
ietf-dots-call-home:source-icmp-type-range	list	32770	4 array	Array
lower-type	uint8	32771	0 unsigned	Number
upper-type	uint8	32772	0 unsigned	Number
ietf-dots-call-home:alt-ch-client	inet: domain-name	32773	3 text string	String
ietf-dots-call-home:alt-ch-client-record	leaf-list inet:ip-address	32774	4 array 3 text string	Array String
ietf-dots-call-home:ttl	uint32	32775	0 unsigned	Number

Table 1: YANG/JSON Mapping Parameters to CBOR

The YANG/JSON mappings to CBOR for "lower-port" and "upper-port" are already defined in Table 5 of [RFC9132].

6.3. YANG Module

This module uses the common YANG types defined in [RFC6991] and the data structure extension defined in [RFC8791].

```
<CODE BEGINS> file "ietf-dots-call-home@2021-12-09.yang"
module ietf-dots-call-home {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-dots-call-home";
  prefix dots-call-home;

  import ietf-inet-types {
    prefix inet;
    reference
      "Section 4 of RFC 6991";
  }
  import ietf-dots-signal-channel {
    prefix dots-signal;
    reference
      "RFC 9132: Distributed Denial-of-Service Open Threat
       Signaling (DOTS) Signal Channel Specification";
  }
  import ietf-yang-structure-ext {
    prefix sx;
    reference
      "RFC 8791: YANG Data Structure Extensions";
  }
}
```

```

}

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

  Author:     Konda, Tirumaleswar Reddy
              <mailto:kondtir@gmail.com>;

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

  Author:     Jon Shallow
              <mailto:ietf-supjps@jpshallow.com>";
description
  "This module contains YANG definitions for the signaling
  messages exchanged between a DOTS client and a DOTS server
  for the Call Home deployment scenario.

  Copyright (c) 2021 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 9066; see
  the RFC itself for full legal notices.";

revision 2021-12-09 {
  description
    "Initial revision.";
  reference
    "RFC 9066: Distributed Denial-of-Service Open Threat
    Signaling (DOTS) Signal Channel Call Home";
}
sx:augment-structure "/dots-signal:dots-signal"
  + "/dots-signal:message-type"
  + "/dots-signal:mitigation-scope"
  + "/dots-signal:scope" {
  description
    "Attack source details.";
  leaf-list source-prefix {
    type inet:ip-prefix;
    description
      "IPv4 or IPv6 prefix identifying the attack source(s).";
  }
  list source-port-range {
    key "lower-port";
    description
      "Port range. When only lower-port is
      present, it represents a single port number.";
    leaf lower-port {
      type inet:port-number;
      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 the lower port number.";
    }
    description
        "Upper port number of the port range.";
}
}
list source-icmp-type-range {
    key "lower-type";
    description
        "ICMP/ICMPv6 type range. When only lower-type is
        present, it represents a single ICMP/ICMPv6 type.

        The address family of the target-prefix is used
        to determine whether ICMP or ICMPv6 is used.";
    leaf lower-type {
        type uint8;
        description
            "Lower ICMP/ICMPv6 type of the ICMP type range.";
        reference
            "RFC 792: Internet Control Message Protocol
            RFC 4443: Internet Control Message Protocol (ICMPv6)
            for the Internet Protocol Version 6 (IPv6)
            Specification.";
    }
    leaf upper-type {
        type uint8;
        must '. >= ../lower-type' {
            error-message
                "The upper ICMP/ICMPv6 type must be greater than
                or equal to the lower ICMP type.";
        }
        description
            "Upper type of the ICMP type range.";
        reference
            "RFC 792: Internet Control Message Protocol
            RFC 4443: Internet Control Message Protocol (ICMPv6)
            for the Internet Protocol Version 6 (IPv6)
            Specification.";
    }
}
}
}
sx:augment-structure "/dots-signal:dots-signal"
    + "/dots-signal:message-type"
    + "/dots-signal:redirection-signal" {
    description
        "Augments the redirection signal to communicate an
        alternate Call Home DOTS client.";
    choice type {
        description
            "Indicates the type of the DOTS session (e.g., base
            DOTS signal channel, DOTS Call Home).";
        case call-home-only {
            description
                "These attributes appear only in a signal Call Home
                channel message from a Call Home DOTS client
                to a Call Home DOTS server.";
            leaf alt-ch-client {
                type inet:domain-name;
                mandatory true;
                description
                    "FQDN of an alternate Call Home DOTS client.

                    This name is also presented as a reference
                    identifier for authentication purposes.";
            }
        }
    }
}

```

```

leaf-list alt-ch-client-record {
    type inet:ip-address;
    description
        "List of IP addresses for the alternate Call
        Home DOTS client.

        If this data node is not present, a Call Home
        DOTS server resolves the alt-ch-client into
        one or more IP addresses.";
}
leaf ttl {
    type uint32;
    units "seconds";
    description
        "The Time To Live (TTL) of the alternate Call Home
        DOTS client.";
    reference
        "Section 4.6 of RFC 9132";
}
}
}
}
}
}
<CODE ENDS>

```

7. IANA Considerations

7.1. DOTS Signal Channel CBOR Mappings Registry

This specification registers the following comprehension-optional parameters (Table 2) in the IANA "DOTS Signal Channel CBOR Key Values" registry [Key-Map].

Parameter Name	CBOR Key Value	CBOR Major Type	Change Controller	Reference
ietf-dots-call-home:source-prefix	32768	4	IESG	RFC 9066
ietf-dots-call-home:source-port-range	32769	4	IESG	RFC 9066
ietf-dots-call-home:source-icmp-type-range	32770	4	IESG	RFC 9066
lower-type	32771	0	IESG	RFC 9066
upper-type	32772	0	IESG	RFC 9066
ietf-dots-call-home:alt-ch-client	32773	3	IESG	RFC 9066
ietf-dots-call-home:alt-ch-client-record	32774	4	IESG	RFC 9066
ietf-dots-call-home:ttl	32775	0	IESG	RFC 9066

Table 2: Assigned DOTS Signal Channel CBOR Key Values

7.2. New DOTS Conflict Cause

Per this document, IANA has assigned a new code from the "DOTS Signal Channel Conflict Cause Codes" registry [Cause].

Code	Label	Description	Reference
4	request-rejected-legitimate-traffic	Mitigation request rejected. This code is returned by the DOTS server to indicate the attack traffic has been classified as legitimate traffic.	RFC 9066

Table 3: Assigned DOTS Signal Channel Conflict Cause Code

7.3. DOTS Signal Call Home YANG Module

Per this document, IANA has registered the following URI in the "ns" subregistry within the "IETF XML Registry" [RFC3688]:

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

Per this document, IANA has registered the following YANG module in the "YANG Module Names" subregistry [RFC6020] within the "YANG Parameters" registry:

name: ietf-dots-call-home
namespace: urn:ietf:params:xml:ns:yang:ietf-dots-call-home
maintained by IANA: N
prefix: dots-call-home
reference: RFC 9066

8. Security Considerations

This document deviates from classic DOTS signal channel usage by having the DOTS server initiate the (D)TLS connection. Security considerations related to the DOTS signal channel discussed in Section 11 of [RFC9132] and (D)TLS early data discussed in Section 7 of [RFC9132] MUST be considered. DOTS agents MUST authenticate each other using (D)TLS before a DOTS signal channel session is considered valid.

The Call Home function enables a Call Home DOTS server to be reachable by only the intended Call Home DOTS client. Appropriate filters (e.g., access control lists) can be installed on the Call Home DOTS server and network between the Call Home DOTS agents so that only communications from a trusted Call Home DOTS client to the Call Home DOTS server are allowed. These filters can be automatically installed by a Call Home DOTS server based on the configured or discovered peer Call Home DOTS client(s).

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 adding the source address to a drop-list after a set number of unsuccessful authentication attempts.

The DOTS signal Call Home channel can be misused by a misbehaving Call Home DOTS client by arbitrarily signaling legitimate traffic as being attack traffic or falsifying mitigation signals so that some sources are disconnected or some traffic is rate-limited. Such misbehaving Call Home DOTS clients may include sources identified by IP addresses that are used for internal use only (that is, these addresses are not visible outside a Call Home DOTS server domain). Absent explicit policy (e.g., the Call Home DOTS client and server are managed by the same administrative entity), such requests should be discarded by the Call Home DOTS server. More generally, Call Home DOTS servers should not blindly trust mitigation requests from Call Home DOTS clients. For example, Call Home DOTS servers could use the attack flow information contained in a mitigation request to enable a full-fledged packet inspection function to inspect all the traffic from the compromised device to the target. They could also redirect the traffic from the potentially compromised device to the target towards a DDoS mitigation system that can scrub the suspicious traffic without blindly blocking all traffic from the indicated attack source to the target. Call Home DOTS servers can also seek the consent of the DOTS server domain administrator to block the traffic from the potentially compromised device to the target (see Section 5.3.1). The means to seek consent are implementation specific.

Call Home DOTS agents may interact with on-path address sharing functions to retrieve an internal IP address / external IP address mapping (Section 5.3.2) identifying an attack source. Blocking access or manipulating the mapping information will complicate DDoS attack mitigation close to an attack source. Additional security considerations are specific to the actual mechanism used to access that mapping (refer, e.g., to Section 4 of [RFC8512] or Section 4 of [RFC8513]).

This document augments YANG data structures that are meant to be used as an abstract representation of DOTS signal channel Call Home messages. As such, the "ietf-dots-call-home" module does not introduce any new vulnerabilities beyond those specified above and in [RFC9132].

9. Privacy Considerations

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

Concretely, the protocol does not leak any new information that can be used to ease surveillance. In particular, the Call Home DOTS server is not required to share information that is local to its network (e.g., internal identifiers of an attack source) with the Call Home DOTS client. Also, the recommended data to be included in Call Home DOTS messages is a subset of the Layer 3 / Layer 4 information that can be learned from the overall traffic flows that exit the Call Home DOTS server domain. Furthermore, Call Home DOTS clients do not publicly reveal attack identification information; that information is encrypted and only shared with an authorized entity in the domain to which the IP address/prefix is assigned, from which an attack was issued.

The DOTS Call Home does not preclude the validation of mitigation

requests received from a Call Home DOTS client. For example, a security service running on the CPE may require an administrator's consent before the CPE acts upon the mitigation request indicated by the Call Home DOTS client. How the consent is obtained is out of scope of this document.

Note that a Call Home DOTS server can seek an administrator's consent, validate the request by inspecting the relevant traffic for attack signatures, or proceed with both courses of action.

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

Moreover, the DOTS Call Home avoids misattribution by appropriately identifying the network to which a suspect attack source belongs (e.g., address sharing issues discussed in Section 5.3.1).

Triggers to send a DOTS mitigation request to a Call Home DOTS server are deployment specific. For example, a Call Home DOTS client may rely on the output of some DDoS detection systems (flow exports or similar functions) deployed within the DOTS client domain to detect potential outbound DDoS attacks or may rely on abuse claims received from remote victim networks. These systems may be misused to track users and infer their activities. Such misuses are not required to achieve the functionality defined in this document (that is, protect the Internet and avoid altering the IP reputation of source networks). It is out of the scope to identify privacy threats specific to given attack detection technology. The reader may refer, for example, to Section 11.8 of [RFC7011].

10. References

10.1. Normative References

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [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>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.

- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.
- [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>>.
- [RFC8791] Bierman, A., Bjorklund, M., and K. Watsen, "YANG Data Structure Extensions", RFC 8791, DOI 10.17487/RFC8791, June 2020, <<https://www.rfc-editor.org/info/rfc8791>>.
- [RFC9132] Boucadair, M., Ed., Shallow, J., and T. Reddy.K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", RFC 9132, DOI 10.17487/RFC9132, September 2021, <<https://www.rfc-editor.org/info/rfc9132>>.

10.2. Informative References

- [Cause] IANA, "DOTS Signal Channel Conflict Cause Codes", <<https://www.iana.org/assignments/dots/>>.
- [DOTS-MULTIHOMING]
Boucadair, M., Reddy, T., and W. Pan, "Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS)", Work in Progress, Internet-Draft, draft-ietf-dots-multihoming-09, 2 December 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-dots-multihoming-09>>.
- [DTLS13] Rescorla, E., Tschofenig, H., and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3", Work in Progress, Internet-Draft, draft-ietf-tls-dtls13-43, 30 April 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-tls-dtls13-43>>.
- [I2NSF-TERMS]
Hares, S., Strassner, J., Lopez, D. R., Xia, L., and H. Birkholz, "Interface to Network Security Functions (I2NSF) Terminology", Work in Progress, Internet-Draft, draft-ietf-i2nsf-terminology-08, 5 July 2019, <<https://datatracker.ietf.org/doc/html/draft-ietf-i2nsf-terminology-08>>.
- [Key-Map] IANA, "DOTS Signal Channel CBOR Key Values", <<https://www.iana.org/assignments/dots/>>.
- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, DOI 10.17487/RFC2663, August 1999, <<https://www.rfc-editor.org/info/rfc2663>>.

- [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>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.
- [RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <<https://www.rfc-editor.org/info/rfc6398>>.
- [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>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [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>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.
- [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>>.
- [RFC8517] Dolson, D., Ed., Snellman, J., Boucadair, M., Ed., and C. Jacquenet, "An Inventory of Transport-Centric Functions Provided by Middleboxes: An Operator Perspective", RFC 8517, DOI 10.17487/RFC8517, February 2019, <<https://www.rfc-editor.org/info/rfc8517>>.
- [RFC8576] Garcia-Morchon, O., Kumar, S., and M. Sethi, "Internet of Things (IoT) Security: State of the Art and Challenges", RFC 8576, DOI 10.17487/RFC8576, April 2019, <<https://www.rfc-editor.org/info/rfc8576>>.
- [RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <<https://www.rfc-editor.org/info/rfc8612>>.
- [RFC8811] Mortensen, A., Ed., Reddy, K., T., Ed., Andreasen, F., Teague, N., and R. Compton, "DDoS Open Threat Signaling (DOTS) Architecture", RFC 8811, DOI 10.17487/RFC8811, August 2020, <<https://www.rfc-editor.org/info/rfc8811>>.
- [RFC8903] Dobbins, R., Migault, D., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use Cases for DDoS Open Threat Signaling", RFC 8903, DOI 10.17487/RFC8903, May 2021, <<https://www.rfc-editor.org/info/rfc8903>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/info/rfc8955>>.
- [RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <<https://www.rfc-editor.org/info/rfc8956>>.
- [RFC8973] Boucadair, M. and T. Reddy, K., "DDoS Open Threat Signaling (DOTS) Agent Discovery", RFC 8973, DOI 10.17487/RFC8973, January 2021, <<https://www.rfc-editor.org/info/rfc8973>>.
- [RS] RSnake, "Slowloris HTTP DoS", <<https://web.archive.org/web/20150315054838/http://ha.ckers.org/slowloris/>>.
- [Sec-by-design] UK Department for Digital, Culture, Media & Sport, "Secure by Design: Improving the cyber security of consumer Internet of Things Report", March 2018, <<https://www.gov.uk/government/publications/secure-by-design-report>>.

Appendix A. Some Home Network Issues

Internet of Things (IoT) devices are becoming more and more prevalent, in particular in home networks. With compute and memory becoming cheaper and cheaper, various types of IoT devices become available in the consumer market at affordable prices. But on the downside, there is a corresponding threat since most of these IoT devices are bought off-the-shelf and most manufacturers haven't considered security in the product design (e.g., [Sec-by-design]). IoT devices deployed in home networks can be easily compromised, they often do not have an easy mechanism to upgrade, and even when

upgradable, IoT manufacturers may cease manufacture and/or discontinue patching vulnerabilities on IoT devices (Sections 5.4 and 5.5 of [RFC8576]). These vulnerable and compromised devices will continue to 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 (Section 3 of [RFC8576]) on victims 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. Such misbehaviors can cause collateral damage that will affect end users, and can also harm the reputation of an Internet Service Provider (ISP) for being a source of attack traffic.

Nowadays, network devices in a home network can offer network security functions (e.g., firewall [RFC4949] or Intrusion Protection System (IPS) service [I2NSF-TERMS] on a home router) to protect the devices connected to the home network from both external and internal attacks. It is natural to seek to provide DDoS defense in these devices as well, and 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 within the ISP's network.

Some of the DDoS attacks like spoofed RST or FIN packets, Slowloris [RS], and Transport Layer Security (TLS) renegotiation are difficult to detect on a home network device without adversely affecting its performance. The reason is that typically home devices such as 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. The reader may refer to Section 2 of [RFC6398] for a brief definition of slow and fast paths.

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 because of the memory and CPU limitations of the home routers. Furthermore, for certain DDoS attacks the logic needed to distinguish legitimate traffic from attack traffic on a per-packet basis is complex. This complexity is because that the packet itself may look "legitimate" and no attack signature can be identified. The anomaly can be identified only after detailed statistical analysis. In addition, network security services in home networks may not be able to detect all types of DDoS attacks using DPI. ISPs offering DDoS mitigation services have a DDoS detection capability that relies upon anomaly detection to identify zero day DDoS attacks and to detect DDoS attacks that cannot be detected using signatures and rate-limit techniques.

ISPs can detect some DDoS attacks originating from a home network (e.g., Section 2.6 of [RFC8517]), but the ISP usually does not have a mechanism to detect which device in the home network is generating the DDoS attack traffic. The primary reason for this is that devices in an IPv4 home network are typically behind a Network Address Translation (NAT) border [RFC2663]. Even in case of an 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 its IPv6 address to evade remediation. A security function on the local home network is better positioned to track the compromised device across IPv6 address (and potentially even MAC address) changes and thus ensure that remediation remains in place across such events.

Appendix B. Disambiguating Base DOTS Signal vs. DOTS Call Home

With the DOTS signal channel Call Home, there is a chance that two DOTS agents can simultaneously establish two DOTS signal channels with different directions (base DOTS signal channel and DOTS signal channel Call Home). Here is one example drawn from the home network. Nevertheless, the outcome of the discussion is not specific to these networks, but applies to any DOTS Call Home scenario.

In the Call Home scenario, the Call Home DOTS server in, for example, the home network can mitigate the DDoS attacks launched by the compromised device in its domain by receiving the mitigation request sent by the Call Home DOTS client in the ISP environment. In addition, the DOTS client in the home network can initiate a mitigation request to the DOTS server in the ISP environment to ask for help when the home network is under a DDoS attack. Such Call Home DOTS server and DOTS client in the home network can co-locate in the same home network element (e.g., the Customer Premises Equipment). In this case, with the same peer at the same time the home network element will have the base DOTS signal channel defined in [RFC9132] and the DOTS signal channel Call Home defined in this specification. Thus, these two signal channels need to be distinguished when they are both supported. Two approaches have been considered for distinguishing the two DOTS signal channels, but only the one that using the dedicated port number has been chosen as the best choice.

By using a dedicated port number for each, these two signal channels can be separated unambiguously and easily. For example, the CPE uses the port number 4646 allocated in [RFC9132] to initiate the basic signal channel to the ISP when it acts as the DOTS client, and uses another port number to initiate the signal channel Call Home. Based on the different port numbers, the ISP can directly decide which kind of procedures should follow immediately after it receives the DOTS messages. This approach just requires two (D)TLS sessions to be established respectively for the basic signal channel and signal channel Call Home.

The other approach is signaling the role of each DOTS agent (e.g., by using the DOTS data channel as depicted in Figure 14). For example, the DOTS agent in the home network first initiates a DOTS data channel to the peer DOTS agent in the ISP environment, at this time the DOTS agent in the home network is the DOTS client and the peer DOTS agent in the ISP environment is the DOTS server. After that, the DOTS agent in the home network retrieves the DOTS Call Home capability of the peer DOTS agent. If the peer supports the DOTS Call Home, the DOTS agent needs to subscribe to the peer to use this extension. Then, the reversal of DOTS role can be recognized as done by both DOTS agents. When the DOTS agent in the ISP environment, which now is the DOTS client, wants to filter the attackers' traffic, it requests the DOTS agent in the home network, which now is the DOTS server, for help.

```
augment /ietf-data:dots-data/ietf-data:capabilities:
  +--ro call-home-support?   boolean
augment /ietf-data:dots-data/ietf-data:dots-client:
  +--rw call-home-enable?    boolean
```

Figure 14: Example of DOTS Data Channel Augmentation

Signaling the role will complicate the DOTS protocols, and this complexity is not required in context where the DOTS Call Home is not required or only when the DOTS Call Home is needed. Besides, the DOTS data channel may not work during attack time. Even if changing the above example from using the DOTS data channel to the DOTS signal channel, the more procedures will still reduce the efficiency. Using

the dedicated port number is much easier and more concise compared to the second approach, and its cost that establishing two (D)TLS sessions is much less. So, using a dedicated port number for the DOTS Call Home is recommended in this specification. The dedicated port number can be configured locally or discovered using means such as [RFC8973].

Acknowledgements

Thanks to Wei Pei, Xia Liang, Roman Danyliw, Dan Wing, Toema Gavrichenkov, Daniel Migault, Sean Turner, and Valery Smyslov for the comments.

Benjamin Kaduk's AD review is valuable. Many thanks to him for the detailed review.

Thanks to Radia Perlman and David Schinazi for the directorate reviews.

Thanks to Ebben Aries for the YANG Doctors review.

Thanks to ric Vyncke, Roman Danyliw, Barry Leiba, Robert Wilton, and Erik Kline for the IESG review.

Contributors

The following individuals have contributed to this document:

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

Email: harsha_joshi@mcafee.com

Wei Pan
Huawei Technologies
China

Email: william.panwei@huawei.com

Authors' Addresses

Tirumaleswar Reddy.K
Akamai
Embassy Golf Link Business Park
Bangalore 560071
Karnataka
India

Email: kondtir@gmail.com

Mohamed Boucadair (editor)
Orange
35000 Rennes
France

Email: mohamed.boucadair@orange.com

Jon Shallow

United Kingdom

Email: supjps-ietf@jpshallow.com