

DOTS
Internet-Draft
Intended status: Standards Track
Expires: December 31, 2021

M. Boucadair
Orange
J. Shallow
June 29, 2021

Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal
Channel Configuration Attributes for Robust Block Transmission
draft-bosh-dots-quick-blocks-03

Abstract

This document specifies new DOTS signal channel configuration parameters that are negotiated between DOTS peers to enable the use of Q-Block1 and Q-Block2 Options. These options enable robust and faster transmission rates for large amounts of data with less packet interchanges as well as supporting faster recovery should any of the blocks get lost in transmission.

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 December 31, 2021.

Copyright Notice

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

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

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

Table of Contents

1. Introduction	2
2. Terminology	3
3. DOTS Attributes for Robust Block Transmission	4
4. DOTS Fast Block Transmission YANG Module	7
4.1. Tree Structure	8
4.2. YANG/JSON Mapping Parameters to CBOR	9
4.3. YANG Module	10
5. IANA Considerations	16
5.1. DOTS Signal Channel CBOR Mappings Registry	16
5.2. DOTS Signal Filtering Control YANG Module	16
6. Security Considerations	17
7. Acknowledgements	17
8. References	17
8.1. Normative References	17
8.2. Informative References	18
Authors' Addresses	19

1. Introduction

The Constrained Application Protocol (CoAP) [RFC7252], although inspired by HTTP, was designed to use UDP instead of TCP. The message layer of CoAP over UDP includes support for reliable delivery, simple congestion control, and flow control. [RFC7959] introduced the CoAP Block1 and Block2 Options to handle data records that cannot fit in a single IP packet, so not having to rely on IP fragmentation and was further updated by [RFC8323] for use over TCP, TLS, and WebSockets.

The CoAP Block1 and Block2 Options work well in environments where there are no or minimal packet losses. These options operate synchronously where each individual block has to be requested and can only ask for (or send) the next block when the request for the previous block has completed. Packet, and hence block transmission rate, is controlled by Round Trip Times (RTTs).

There is a requirement for these blocks of data to be transmitted at higher rates under network conditions where there may be asymmetrical transient packet loss (i.e., responses may get dropped). An example is when a network is subject to a Distributed Denial of Service (DDoS) attack and there is a need for DDoS mitigation agents relying upon CoAP to communicate with each other (e.g., [I-D.ietf-dots-telemetry]). As a reminder, [RFC7959] recommends the use of Confirmable (CON) responses to handle potential packet loss.

However, such a recommendation does not work with a flooded pipe DDoS situation as the returning ACK packets may not get through.

The block-wise transfer specified in [RFC7959] covers the general case, but falls short in situations where packet loss is highly asymmetrical. The mechanism specified in [I-D.ietf-core-new-block] provides roughly similar features to the Block1/Block2 Options. It provides additional properties that are tailored towards the intended DOTS transmission. Concretely, [I-D.ietf-core-new-block] primarily targets applications such as DDoS Open Threat Signaling (DOTS) that can't use Confirmable (CON) responses to handle potential packet loss and that support application-specific mechanisms to assess whether the remote peer is able to handle the messages sent by a CoAP endpoint (e.g., DOTS heartbeats in Section 4.7 of [I-D.ietf-dots-rfc8782-bis]).

[I-D.ietf-core-new-block] includes guards to prevent a CoAP agent from overloading the network by adopting an aggressive sending rate. These guards are followed in addition to the existing CoAP congestion control as specified in Section 4.7 of [RFC7252]. Table 1 additional CoAP attributes that are used for the guards.

Parameter Name	Default Value
MAX_PAYLOADS	10
NON_MAX_RETRANSMIT	4
NON_TIMEOUT	2 s
NON_RECEIVE_TIMEOUT	4 s
NON_PROBING_WAIT	247 s
NON_PARTIAL_TIMEOUT	247 s

Table 1: Congestion Control Parameters

PROBING_RATE and other transmission parameters are negotiated between DOTS peers as discussed in Section 4.5.2 of [I-D.ietf-dots-rfc8782-bis]. Nevertheless, the attributes listed in Table 1 are not supported. This document defines new DOTS signal channel attributes that are meant to customize the configuration of robust block transmission in a DOTS context.

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.

Readers should be familiar with the terms and concepts defined in [RFC7252] and [RFC8612].

The terms "payload" and "body" are defined in [RFC7959]. The term "payload" is thus used for the content of a single CoAP message (i.e., a single block being transferred), while the term "body" is used for the entire resource representation that is being transferred in a block-wise fashion.

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]. Specific terms are used for any statement that applies to either protocol alone.

3. DOTS Attributes for Robust Block Transmission

Section 6.2 of [I-D.ietf-core-new-block] defines the following attributes that are used for congestion control purposes:

MAX_PAYLOADS: is the maximum number of payloads that can be transmitted at any one time.

NON_MAX_RETRANSMIT: is the maximum number of times a request for the retransmission of missing payloads can occur without a response from the remote peer. By default, NON_MAX_RETRANSMIT has the same value as MAX_RETRANSMIT (Section 4.8 of [RFC7252]).

NON_TIMEOUT: is the maximum period of delay between sending sets of MAX_PAYLOADS payloads for the same body. NON_TIMEOUT has the same value as ACK_TIMEOUT (Section 4.8 of [RFC7252]).

NON_RECEIVE_TIMEOUT: is the maximum time to wait for a missing payload before requesting retransmission. By default, NON_RECEIVE_TIMEOUT has a value of twice NON_TIMEOUT.

NON_PROBING_WAIT: is used to limit the potential wait needed calculated when using PROBING_WAIT. By default, NON_PROBING_WAIT has the same value as EXCHANGE_LIFETIME (Section 4.8.2 of [RFC7252]).

NON_PARTIAL_TIMEOUT: is used for expiring partially received bodies. By default, NON_PARTIAL_TIMEOUT has the same value as EXCHANGE_LIFETIME (Section 4.8.2 of [RFC7252]).

These attributes are used together with PROBING_RATE parameter which in CoAP indicates the average data rate that must not be exceeded by a CoAP endpoint in sending to a peer endpoint that does not respond. The single body of blocks will be subjected to PROBING_RATE (Section 4.7 of [RFC7252]), not the individual packets. If the wait time between sending bodies that are not being responded to calculated using on PROBING_RATE exceeds NON_PROBING_WAIT, then the gap time is limited to NON_PROBING_WAIT.

This document augments the "ietf-dots-signal-channel" (dots-signal) DOTS signal YANG module defined in [I-D.ietf-dots-rfc8782-bis] with these additional attributes that can be negotiated between DOTS peers to enable robust and faster transmission:

max-payloads: This attribute echoes the MAX_PAYLOADS parameter in [I-D.ietf-core-new-block].

This is an optional attribute.

For the sake of more flexible configuration, this document defines also the following attributes:

non-max-retransmit: This attribute echoes the NON_MAX_RETRANSMIT parameter in [I-D.ietf-core-new-block]. The default value of this attribute is 'max-retransmit'. Note that DOTS uses a default value of '3' instead of '4' used for the generic CoAP use (Section 4.5.2 of [I-D.ietf-dots-rfc8782-bis]) for max-transmit.

This is an optional attribute.

non-timeout: This attribute echoes the NON_TIMEOUT parameter in [I-D.ietf-core-new-block]. The default value of this attribute is 'ack-timeout'.

This is an optional attribute.

non-probing-wait: This attribute echoes the NON_PROBING_WAIT parameter in [I-D.ietf-core-new-block]. The default value of this attribute is 247s.

This is an optional attribute.

non-partial-timeout: This attribute echoes the NON_PARTIAL_TIMEOUT parameter in [I-D.ietf-core-new-block]. The default value of this attribute is 274s.

This is an optional attribute.

An example of PUT message to convey the configuration parameters for the DOTS signal channel is depicted in Figure 1. In this example, the 'max-payloads' is set to '15' when no mitigation is active, while it is set to '10' when a mitigation is active. The same value is used for both 'non-max-retransmit' and 'non-timeout' in idle and mitigation times.

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

{
  "ietf-dots-signal-channel:signal-config": {
    "mitigating-config": {
      "heartbeat-interval": {
        "current-value": 30
      },
      "missing-hb-allowed": {
        "current-value": 15
      },
      "probing-rate": {
        "current-value": 15
      },
      "max-retransmit": {
        "current-value": 3
      },
      "ack-timeout": {
        "current-value-decimal": "2.00"
      },
      "ack-random-factor": {
        "current-value-decimal": "1.50"
      },
      "ietf-dots-robust-trans:max-payloads": {
        "current-value": 10
      },
      "ietf-dots-robust-trans:non-max-retransmit": {
        "current-value": 3
      },
      "ietf-dots-robust-trans:non-timeout": {
```

```

        "current-value-decimal": "2.00"
      },
      "ietf-dots-robust-trans:non-probing-wait": {
        "current-value-decimal": "247.00"
      },
      "ietf-dots-robust-trans:non-partial-wait": {
        "current-value-decimal": "247.00"
      }
    },
    "idle-config": {
      "heartbeat-interval": {
        "current-value": 0
      },
      "max-retransmit": {
        "current-value": 3
      },
      "ack-timeout": {
        "current-value-decimal": "2.00"
      },
      "ack-random-factor": {
        "current-value-decimal": "1.50"
      },
      "ietf-dots-robust-trans:max-payloads": {
        "current-value": 15
      },
      "ietf-dots-robust-trans:non-max-retransmit": {
        "current-value": 3
      },
      "ietf-dots-robust-trans:non-timeout": {
        "current-value-decimal": "2.00"
      },
      "ietf-dots-robust-trans:non-probing-wait": {
        "current-value-decimal": "247.00"
      },
      "ietf-dots-robust-trans:non-partial-wait": {
        "current-value-decimal": "247.00"
      }
    }
  }
}

```

Figure 1: Example of PUT to Convey the Configuration Parameters

4. DOTS Fast Block Transmission YANG Module

4.1. Tree Structure

This document defines the YANG module "ietf-dots-robust-trans" (Section 4), which has the following tree structure:

module: ietf-dots-robust-trans

```
augment-structure /dots-signal:dots-signal/dots-signal:message-type
                  /dots-signal:signal-config
                  /dots-signal:mitigating-config:
```

```
+-- max-payloads
|   +-- (direction)?
|   |   +--:(server-to-client-only)
|   |   |   +-- max-value?      uint16
|   |   |   +-- min-value?      uint16
|   |   +-- current-value?      uint16
|   +-- non-max-retransmit
|   |   +-- (direction)?
|   |   |   +--:(server-to-client-only)
|   |   |   |   +-- max-value?      uint16
|   |   |   |   +-- min-value?      uint16
|   |   +-- current-value?      uint16
|   +-- non-timeout
|   |   +-- (direction)?
|   |   |   +--:(server-to-client-only)
|   |   |   |   +-- max-value-decimal? decimal64
|   |   |   |   +-- min-value-decimal? decimal64
|   |   +-- current-value-decimal? decimal64
|   +-- non-probing-wait
|   |   +-- (direction)?
|   |   |   +--:(server-to-client-only)
|   |   |   |   +-- max-value-decimal? decimal64
|   |   |   |   +-- min-value-decimal? decimal64
|   |   +-- current-value-decimal? decimal64
|   +-- non-partial-wait:
|   |   +-- (direction)?
|   |   |   +--:(server-to-client-only)
|   |   |   |   +-- max-value-decimal? decimal64
|   |   |   |   +-- min-value-decimal? decimal64
|   |   +-- current-value-decimal? decimal64
```

```
augment-structure /dots-signal:dots-signal/dots-signal:message-type
                  /dots-signal:signal-config/dots-signal:idle-config:
```

```
+-- max-payloads
|   +-- (direction)?
|   |   +--:(server-to-client-only)
|   |   |   +-- max-value?      uint16
|   |   |   +-- min-value?      uint16
```



```

|   +-- current-value?      uint16
+-- non-max-retransmit
|   +-- (direction)?
|       +--:(server-to-client-only)
|           +-- max-value?    uint16
|           +-- min-value?    uint16
|   +-- current-value?      uint16
+-- non-timeout
|   +-- (direction)?
|       +--:(server-to-client-only)
|           +-- max-value-decimal? decimal64
|           +-- min-value-decimal? decimal64
|   +-- current-value-decimal? decimal64
+-- non-probing-wait
|   +-- (direction)?
|       +--:(server-to-client-only)
|           +-- max-value-decimal? decimal64
|           +-- min-value-decimal? decimal64
|   +-- current-value-decimal? decimal64
+-- non-partial-wait:
|   +-- (direction)?
|       +--:(server-to-client-only)
|           +-- max-value-decimal? decimal64
|           +-- min-value-decimal? decimal64
|   +-- current-value-decimal? decimal64

```

4.2. YANG/JSON Mapping Parameters to CBOR

The YANG/JSON mapping parameters to CBOR are listed in Table 2.

- o Note: Implementers must check that the mapping output provided by their YANG-to-CBOR encoding schemes is aligned with the content of Table 2.

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
ietf-dots-robust-trans:max-payloads	container	TBA1	5 map	Object
ietf-dots-robust-trans:non-max-retransmit	container	TBA2	5 map	Object
ietf-dots-robust-trans:non-timeout	container	TBA3	5 map	Object
ietf-dots-robust-trans:non-probing-wait	container	TBA4	5 map	Object
ietf-dots-robust-trans:non-partial-wait	container	TBA5	5 map	Object

Table 2: YANG/JSON Mapping Parameters to CBOR

4.3. YANG Module

This module uses the data structure extension defined in [RFC8791].

```
<CODE BEGINS> file "ietf-dots-robust-trans@2020-05-04.yang"
module ietf-dots-robust-trans {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-dots-robust-trans";
  prefix dots-robust;

  import ietf-dots-signal-channel {
    prefix dots-signal;
    reference
      "RFC YYYY: 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: Mohamed Boucadair

<<mailto:mohamed.boucadair@orange.com>>;

Author: Jon Shallow

<<mailto:ietf-suppjps@jpshallow.com>>;

description

"This module contains YANG definitions for the configuration of parameters that can be negotiated between a DOTS client and a DOTS server for robust block transmission.

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 XXXX; see the RFC itself for full legal notices.";

revision 2021-05-04 {

description

"Initial revision.";

reference

"RFC XXXX: Distributed Denial-of-Service Open Threat Signaling (DOTS) Configuration Attributes for Robust Block Transmission";

}

grouping robust-transmission-attributes {

description

"A set of DOTS signal channel session configuration that are negotiated between DOTS agents when making use of Q-Block1 and Q-Block2 Options.";

container max-payloads {

description

"Indicates the maximum number of payloads that can be transmitted at any one time.";

choice direction {

description

"Indicates the communication direction in which the

```
        data nodes can be included.";
    case server-to-client-only {
        description
            "These data nodes appear only in a mitigation message
            sent from the server to the client.";
        leaf max-value {
            type uint16;
            description
                "Maximum acceptable max-payloads value.";
        }
        leaf min-value {
            type uint16;
            description
                "Minimum acceptable max-payloads value.";
        }
    }
}
leaf current-value {
    type uint16;
    default "10";
    description
        "Current max-payloads value.";
}
}
container non-max-retransmit {
    description
        "Indicates the the maximum number of times a
        request for the retransmission of missings payloads
        can occur without a response from the remote peer.";
    leaf max-value {
        type uint16;
        config false;
        description
            "Maximum acceptable non-max-retransmit value.";
    }
    leaf min-value {
        type uint16;
        config false;
        description
            "Minimum acceptable non-max-retransmit value.";
    }
    leaf current-value {
        type uint16;
        default "3";
        description
            "Current non-max-retransmit value.";
    }
}
}
```

```
container non-timeout {
  description
    "Indicates the maximum period of delay between
    sending sets of MAX_PAYLOADS payloads for the same
    body. By default, this parameter has the same value
    as ACK_TIMEOUT.";
  choice direction {
    description
      "Indicates the communication direction in which the
      data nodes can be included.";
    case server-to-client-only {
      description
        "These data nodes appear only in a mitigation message
        sent from the server to the client.";
      leaf max-value-decimal {
        type decimal64 {
          fraction-digits 2;
        }
        units "seconds";
        description
          "Maximum ack-timeout value.";
      }
      leaf min-value-decimal {
        type decimal64 {
          fraction-digits 2;
        }
        units "seconds";
        description
          "Minimum ack-timeout value.";
      }
    }
  }
  leaf current-value-decimal {
    type decimal64 {
      fraction-digits 2;
    }
    units "seconds";
    default "2";
    description
      "Current ack-timeout value.";
  }
}

container non-probing-wait {
  description
    "Is used to limit the potential wait needed calculated
    when using probing-rate.";
  choice direction {
    description
```

```
    "Indicates the communication direction in which the
      data nodes can be included.";
  case server-to-client-only {
    description
      "These data nodes appear only in a mitigation message
        sent from the server to the client.";
    leaf max-value-decimal {
      type decimal64 {
        fraction-digits 2;
      }
      units "seconds";
      description
        "Maximum non-probing-wait value.";
    }
    leaf min-value-decimal {
      type decimal64 {
        fraction-digits 2;
      }
      units "seconds";
      description
        "Minimum non-probing-wait value.";
    }
  }
}
leaf current-value-decimal {
  type decimal64 {
    fraction-digits 2;
  }
  units "seconds";
  default "247";
  description
    "Current non-probing-wait value.";
}
}
container non-partial-wait {
  description
    "Is used for expiring partially received bodies.";
  choice direction {
    description
      "Indicates the communication direction in which the
        data nodes can be included.";
    case server-to-client-only {
      description
        "These data nodes appear only in a mitigation message
          sent from the server to the client.";
      leaf max-value-decimal {
        type decimal64 {
          fraction-digits 2;
        }

```

```
    }
    units "seconds";
    description
        "Maximum non-partial-wait value.";
    }
    leaf min-value-decimal {
        type decimal64 {
            fraction-digits 2;
        }
        units "seconds";
        description
            "Minimum non-partial-wait value.";
    }
}
}
leaf current-value-decimal {
    type decimal64 {
        fraction-digits 2;
    }
    units "seconds";
    default "247";
    description
        "Current non-partial-wait value.";
}
}
}

sx:augment-structure "/dots-signal:dots-signal"
    + "/dots-signal:message-type"
    + "/dots-signal:signal-config"
    + "/dots-signal:mitigating-config" {
    description
        "Indicates DOTS configuration parameters to use for
        robust transmission when a mitigation is active.";
    uses robust-transmission-attributes;
}
sx:augment-structure "/dots-signal:dots-signal"
    + "/dots-signal:message-type"
    + "/dots-signal:signal-config"
    + "/dots-signal:idle-config" {
    description
        "Indicates DOTS configuration parameters to use for
        robust transmission when no mitigation is active.";
    uses robust-transmission-attributes;
}
}
<CODE ENDS>
```

5. IANA Considerations

5.1. DOTS Signal Channel CBOR Mappings Registry

This specification registers the following parameters in the IANA "DOTS Signal Channel CBOR Key Values" registry [Key-Map].

- o Note to the RFC Editor: Please replace TBA1/TBA2/TBA3 with the CBOR keys that are assigned from the 128-255 range. Please update Table 2 accordingly.

Parameter Name	CBOR Key Value	CBOR Major Type	Change Controller	Specification Document(s)
ietf-dots-robust-trans: max-payloads	TBA1	5	IESG	[RFCXXXX]
ietf-dots-robust-trans: non-max-retransmit	TBA2	5	IESG	[RFCXXXX]
ietf-dots-robust-trans: non-timeout	TBA3	5	IESG	[RFCXXXX]
ietf-dots-robust-trans: non-probing-wait	TBA4	5	IESG	[RFCXXXX]
ietf-dots-robust-trans: non-partial-wait	TBA5	5	IESG	[RFCXXXX]

5.2. DOTS Signal Filtering Control YANG Module

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

URI: urn:ietf:params:xml:ns:yang:ietf-dots-robust-trans
 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" subregistry [RFC6020] within the "YANG Parameters" registry.

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

6. Security Considerations

The security considerations for the DOTS signal channel protocol are discussed in Section 11 of [I-D.ietf-dots-rfc8782-bis].

CoAP-specific security considerations are discussed in Section 11 of [I-D.ietf-core-new-block].

This document defines YANG data structures that are meant to be used as an abstract representation in DOTS signal channel messages. As such, the "ietf-dots-robust-trans" module does not introduce any new vulnerabilities beyond those specified above.

7. Acknowledgements

TBC

8. References

8.1. Normative References

- [I-D.ietf-core-new-block]
Boucadair, M. and J. Shallow, "Constrained Application Protocol (CoAP) Block-Wise Transfer Options for Faster Transmission", draft-ietf-core-new-block-11 (work in progress), April 2021.
- [I-D.ietf-dots-rfc8782-bis]
Boucadair, M., Shallow, J., and T. Reddy.K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-rfc8782-bis-06 (work in progress), March 2021.
- [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>>.

- [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>>.
- [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>>.
- [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>>.
- [RFC7959] Bormann, C. and Z. Shelby, Ed., "Block-Wise Transfers in the Constrained Application Protocol (CoAP)", RFC 7959, DOI 10.17487/RFC7959, August 2016, <<https://www.rfc-editor.org/info/rfc7959>>.
- [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>>.
- [RFC8323] Bormann, C., Lemay, S., Tschofenig, H., Hartke, K., Silverajan, B., and B. Raymor, Ed., "CoAP (Constrained Application Protocol) over TCP, TLS, and WebSockets", RFC 8323, DOI 10.17487/RFC8323, February 2018, <<https://www.rfc-editor.org/info/rfc8323>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-dots-telemetry] Boucadair, M., Reddy, T., Doron, E., Chen, M., and J. Shallow, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Telemetry", draft-ietf-dots-telemetry-15 (work in progress), December 2020.
- [Key-Map] IANA, "DOTS Signal Channel CBOR Key Values", <<https://www.iana.org/assignments/dots/dots.xhtml#dots-signal-channel-cbor-key-values>>.

- [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>>.
- [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>>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Jon Shallow
United Kingdom

Email: supjps-ietf@jpshallow.com

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

M. Boucadair
Orange
T. Reddy.K
Akamai
W. Pan
Huawei Technologies
26 April 2022

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

Abstract

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 October 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Requirements Language	4
3. Terminology	4
4. Multi-Homing Scenarios	5
4.1. Multi-Homed Residential Single CPE	5
4.2. Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs	6
4.3. Multi-homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	7
4.4. Multi-homed Enterprise with the Same ISP	7
5. DOTS Multi-homing Deployment Considerations	8
5.1. Residential CPE	8
5.2. Multi-Homed Enterprise: Single CPE, Multiple Upstream ISPs	10
5.3. Multi-Homed Enterprise: Multiple CPEs, Multiple Upstream ISPs	12
5.4. Multi-Homed Enterprise: Single ISP	13
6. Security Considerations	14
7. IANA Considerations	14
8. Acknowledgements	14
9. References	15
9.1. Normative References	15
9.2. Informative References	15
Authors' Addresses	16

1. Introduction

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

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

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




Figure 1: Basic DOTS Architecture

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

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

1. Complete the base DOTS architecture with multi-homing specifics. Those specifics need to be taken into account because:
 - * Sending a DOTS mitigation request to an arbitrary DOTS server will not necessarily help in mitigating a DDoS attack.
 - * Randomly replicating all DOTS mitigation requests among all available DOTS servers is suboptimal.
 - * Sequentially contacting DOTS servers may increase the delay before a mitigation plan is enforced.
2. Identify DOTS deployment schemes in a multi-homing context, where DOTS services can be offered by all or a subset of upstream providers.

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

- * Select the appropriate DOTS server(s).
- * Identify cases where anycast is not recommended for DOTS.

This document adopts the following methodology:

- * Identify and extract viable deployment candidates from [RFC8903].
- * Augment the description with multi-homing technicalities, e.g.,
 - One vs. multiple upstream network providers
 - One vs. multiple interconnect routers
 - Provider-Independent (PI) vs. Provider-Aggregatable (PA) IP addresses
- * Describe the recommended behavior of DOTS clients and client-domain DOTS gateways for each case.

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

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology

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

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

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

IP indifferently refers to IPv4 or IPv6.

4. Multi-Homing Scenarios

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

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

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

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

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

4.1. Multi-Homed Residential Single CPE

The scenario shown in Figure 2 is characterized as follows:

- * The home network is connected to the Internet using one single CPE.
- * The CPE is connected to multiple provisioning domains (i.e., both fixed and mobile networks). Provisioning domain (PvD) is explained in [RFC7556].

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

- * Each of these provisioning domains assigns IP addresses/prefixes to the CPE and provides additional configuration information such as a list of DNS servers, DNS suffixes associated with the network, default gateway address, and DOTS server's name [RFC8973]. These addresses/prefixes are assumed to be Provider-Aggregatable (PA).
- * Because of ingress filtering, packets forwarded by the CPE towards a given provisioning domain must be sent with a source IP address that was assigned by that domain [RFC8043].




Figure 2: Typical Multi-homed Residential CPE

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

The scenario shown in Figure 3 is characterized as follows:

- * The enterprise network is connected to the Internet using a single router.
- * That router is connected to multiple provisioning domains managed by distinct administrative entities.

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

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

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




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

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

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

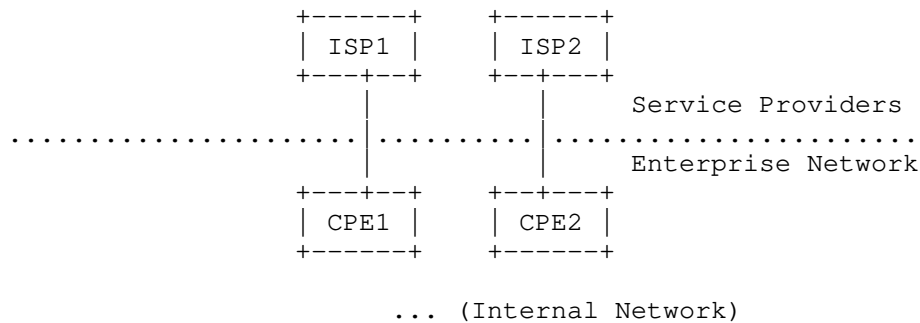


Figure 4: Multi-homed Enterprise Network (Multiple CPEs, Multiple ISPs)

4.4. Multi-homed Enterprise with the Same ISP

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

5. DOTS Multi-homing Deployment Considerations

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

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

Table 1: Sample Deployment Cases

These deployment schemes are further discussed in the following subsections.

5.1. Residential CPE

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




Figure 5: DOTS Associations for a Multihomed Residential CPE

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

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

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

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

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

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

- * One or more DOTS clients are enabled in hosts located in the internal network.
- * A client-domain DOTS gateway is enabled to aggregate and then relay the requests towards upstream DOTS servers.




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

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

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

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

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




Figure 7: Multiple DOTS Clients, Multiple DOTS Servers

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

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

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




Figure 8: Single Homed DOTS Clients

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

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

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

An alternative deployment scheme is shown in Figure 9:

- * DOTS clients are enabled in hosts located in the internal network.
- * A client-domain DOTS gateway is enabled in each CPE (CPE1 and CPE2 per Table 1).

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

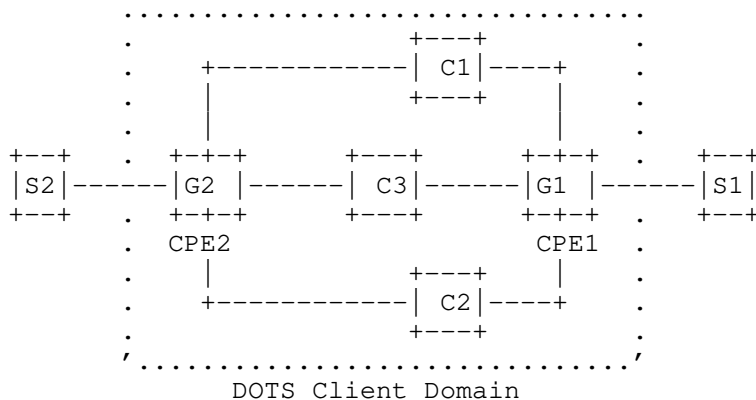


Figure 9: Multiple DOTS Clients, Multiple DOTS Gateways, Multiple DOTS Servers

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

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

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

5.4. Multi-Homed Enterprise: Single ISP

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

- * The same DOTS server for all network attachments.
- * Distinct DOTS servers for each network attachment. These DOTS servers need to coordinate when a mitigation action is received from the enterprise network.

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

6. Security Considerations

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

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

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

7. IANA Considerations

This document does not require any action from IANA.

8. Acknowledgements

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

Thanks to Kirill Kasavchenko for the comments.

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

Many thanks to Roman Danyliw for the careful AD review.

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

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8811] Mortensen, A., Ed., Reddy, K., T., Ed., Andreasen, F., Teague, N., and R. Compton, "DDoS Open Threat Signaling (DOTS) Architecture", RFC 8811, DOI 10.17487/RFC8811, August 2020, <<https://www.rfc-editor.org/info/rfc8811>>.

9.2. Informative References

- [RFC3582] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, DOI 10.17487/RFC3582, August 2003, <<https://www.rfc-editor.org/info/rfc3582>>.
- [RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, DOI 10.17487/RFC4116, July 2005, <<https://www.rfc-editor.org/info/rfc4116>>.
- [RFC4218] Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions", RFC 4218, DOI 10.17487/RFC4218, October 2005, <<https://www.rfc-editor.org/info/rfc4218>>.
- [RFC4732] Handley, M., Ed., Rescorla, E., Ed., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, DOI 10.17487/RFC4732, December 2006, <<https://www.rfc-editor.org/info/rfc4732>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <<https://www.rfc-editor.org/info/rfc7556>>.

- [RFC8043] Sarikaya, B. and M. Boucadair, "Source-Address-Dependent Routing and Source Address Selection for IPv6 Hosts: Overview of the Problem Space", RFC 8043, DOI 10.17487/RFC8043, January 2017, <<https://www.rfc-editor.org/info/rfc8043>>.
- [RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <<https://www.rfc-editor.org/info/rfc8612>>.
- [RFC8783] Boucadair, M., Ed. and T. Reddy.K, Ed., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", RFC 8783, DOI 10.17487/RFC8783, May 2020, <<https://www.rfc-editor.org/info/rfc8783>>.
- [RFC8803] Bonaventure, O., Ed., Boucadair, M., Ed., Gundavelli, S., Seo, S., and B. Hesmans, "0-RTT TCP Convert Protocol", RFC 8803, DOI 10.17487/RFC8803, July 2020, <<https://www.rfc-editor.org/info/rfc8803>>.
- [RFC8903] Dobbins, R., Migault, D., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use Cases for DDoS Open Threat Signaling", RFC 8903, DOI 10.17487/RFC8903, May 2021, <<https://www.rfc-editor.org/info/rfc8903>>.
- [RFC8973] Boucadair, M. and T. Reddy.K, "DDoS Open Threat Signaling (DOTS) Agent Discovery", RFC 8973, DOI 10.17487/RFC8973, January 2021, <<https://www.rfc-editor.org/info/rfc8973>>.
- [RFC9132] Boucadair, M., Ed., Shallow, J., and T. Reddy.K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", RFC 9132, DOI 10.17487/RFC9132, September 2021, <<https://www.rfc-editor.org/info/rfc9132>>.
- [TS.24008] 3GPP, "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3 (Release 16)", December 2019, <<http://www.3gpp.org/DynaReport/24008.htm>>.

Authors' Addresses

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

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

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

DOTS
Internet-Draft
Intended status: Standards Track
Expires: November 27, 2021

T. Reddy
McAfee
M. Boucadair, Ed.
Orange
J. Shallow
May 26, 2021

Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal
Channel Call Home
draft-ietf-dots-signal-call-home-14

Abstract

This document specifies the 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.

Editorial Note (To be removed by RFC Editor)

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

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

Please update this statement with the RFC number to be assigned to the following documents:

- o "RFC YYYY: Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification" (used to be I-D.ietf-dots-rfc8782-bis)

Please update TBD/TBA statements with the assignments made by IANA to DOTS Signal Channel Call Home.

Also, please update the "revision" date of the YANG module.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 27, 2021.

Copyright Notice

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

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

Table of Contents

1. Introduction	3
2. Terminology	6
3. Applicability Scope	7
4. Co-existence of Base DOTS Signal Channel and DOTS Call Home .	8
5. DOTS Signal Channel Call Home	12
5.1. Procedure	12
5.2. DOTS Signal Channel Variations	14
5.2.1. Heartbeat Mechanism	14
5.2.2. Redirected Signaling	15

5.3.	DOTS Signal Channel Extension	16
5.3.1.	Mitigation Request	16
5.3.2.	Address Sharing Considerations	20
6.	DOTS Signal Call Home YANG Module	23
6.1.	Tree Structure	23
6.2.	YANG/JSON Mapping Parameters to CBOR	24
6.3.	YANG Module	25
7.	IANA Considerations	29
7.1.	DOTS Signal Channel CBOR Mappings Registry	29
7.2.	New DOTS Conflict Cause	30
7.3.	DOTS Signal Call Home YANG Module	31
8.	Security Considerations	31
9.	Privacy Considerations	33
10.	Contributors	34
11.	Acknowledgements	34
12.	References	34
12.1.	Normative References	35
12.2.	Informative References	36
Appendix A.	Some Home Network Issues	39
Appendix B.	Disambiguating Base DOTS Signal vs. DOTS Call Home	41
Authors' Addresses	42

1. Introduction

The Distributed Denial-of-Service Open Threat Signaling (DOTS) signal channel protocol [I-D.ietf-dots-rfc8782-bis] 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 [I-D.ietf-dots-use-cases].

However, [I-D.ietf-dots-rfc8782-bis] only covers how to mitigate when being attacked (i.e., protect 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 which 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 is about 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 to 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). DOTS Call Home agents may (or not) be co-located with DOTS agents that are compliant with [I-D.ietf-dots-rfc8782-bis] (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, 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 by the user. Unlike classic DOTS deployments [I-D.ietf-dots-use-cases], a Call Home DOTS server maintains a single DOTS signal channel session for each DOTS-capable upstream provisioning domain [I-D.ietf-dots-multihoming].

For instance, the Call Home DOTS server in the home network initiates the signal channel Call Home in 'idle' time and then 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 [I-D.ietf-dots-rfc8782-bis].

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]. 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.

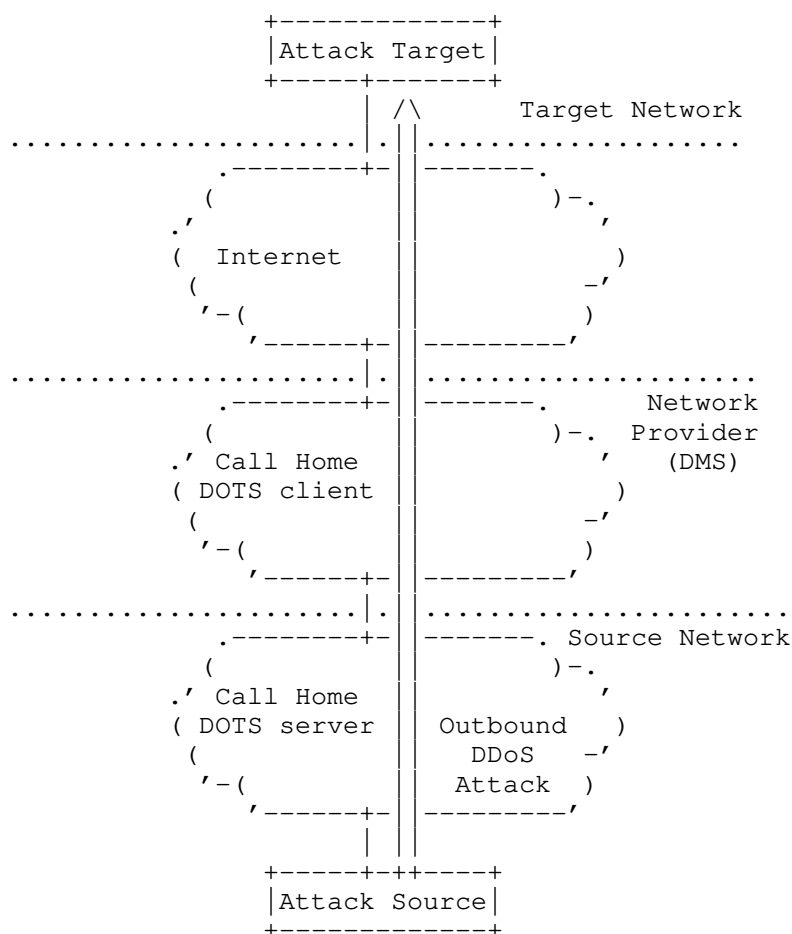


Figure 3: DOTS Signal Channel Call Home Reference Architecture

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 decision is implementation-specific.

For some mitigations, a feedback may be required from an administrator to confirm a filtering action. Means to seek for an administrator's consent are deployment-specific. Indeed, a variety of implementation options can be considered as a function of the Call Home DOTS deployment: 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 be 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 means to notify the user when a new device is detected in order to seek a confirmation whether access should be granted or not 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. Co-existence of Base DOTS Signal Channel and DOTS Call Home

The DOTS signal channel Call Home does not require nor preclude the activation of the base DOTS signal channel [I-D.ietf-dots-rfc8782-bis]. 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 Base DOTS Signal Channel and Call Home (Same DMS Provider)

Note that a DMS provider may not be on the default forwarding path of an 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.

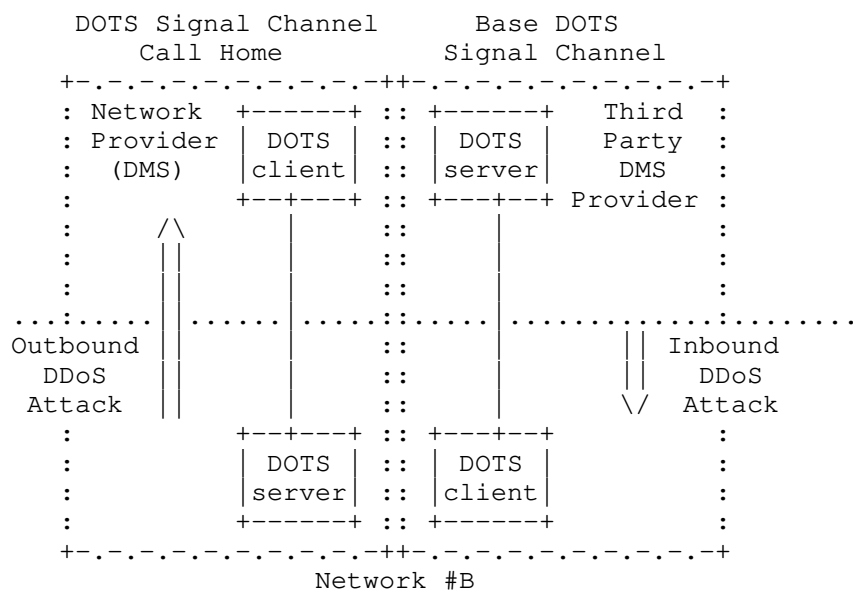


Figure 5: Activation of Base DOTS Signal Channel and Call Home (Distinct DMS Providers)

Figures 6 and 7 depict examples where the same node embeds both base DOTS and DOTS Call Home 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 DOTS Call Home agents is specific to each domain.

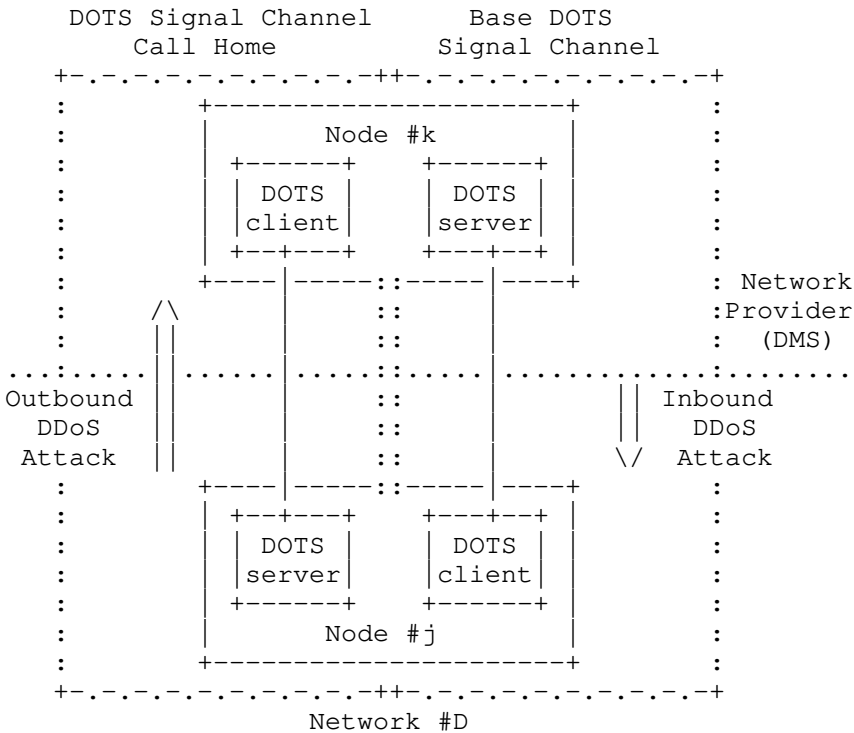


Figure 7: Another Example where the Same Node Embeds both a DOTS Client and a Call Home DOTS Server

Appendix B elaborates on the considerations to unambiguously distinguish DOTS messages which 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 DOTS client-initiated DOTS signal channel protocol [I-D.ietf-dots-rfc8782-bis]. 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 (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 [I-D.ietf-dots-rfc8782-bis] 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 [I-D.ietf-dots-rfc8782-bis]. 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 [I-D.ietf-dots-rfc8782-bis]). The Call Home DOTS server adjusts the 'heartbeat-interval' to accommodate binding timers used by on-path NATs and firewalls. Heartbeats will be then 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 missing heartbeats allowed threshold ('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 [I-D.ietf-dots-rfc8782-bis] (i.e., a 5.03 response that conveys an alternate DOTS server's FQDN or alternate DOTS server's IP address(es)). A Call Home DOTS client MUST silently discard such 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 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 that 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 [I-D.ietf-dots-rfc8782-bis]. 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:t1": 600
  }
}
```

Figure 9: Example of a PUT Request for Redirected Signaling

5.3. DOTS Signal Channel Extension

5.3.1. Mitigation Request

This specification extends the mitigation request defined in Section 4.4.1 of [I-D.ietf-dots-rfc8782-bis] to convey the attack source information (e.g., source prefixes, 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 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 as 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 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 routeable 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 [I-D.ietf-dots-rfc8782-bis].

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.

The 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 a consent from the Call Home DOTS server domain administrator is required, the Call Home DOTS server replies with 2.01 (Created) and 'status' code set to 1 (attack-mitigation-in-progress). Then, the mechanisms defined in Section 4.4.2 of [I-D.ietf-dots-rfc8782-bis] are followed by the DOTS agents to update the mitigation status. Particularly, 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 [I-D.ietf-dots-rfc8782-bis]) 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 [I-D.ietf-dots-rfc8782-bis].

The DOTS agents follow the same procedures specified in [I-D.ietf-dots-rfc8782-bis] 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 [RFC8512] or [RFC8513], to retrieve the internal IP address and port number which 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.

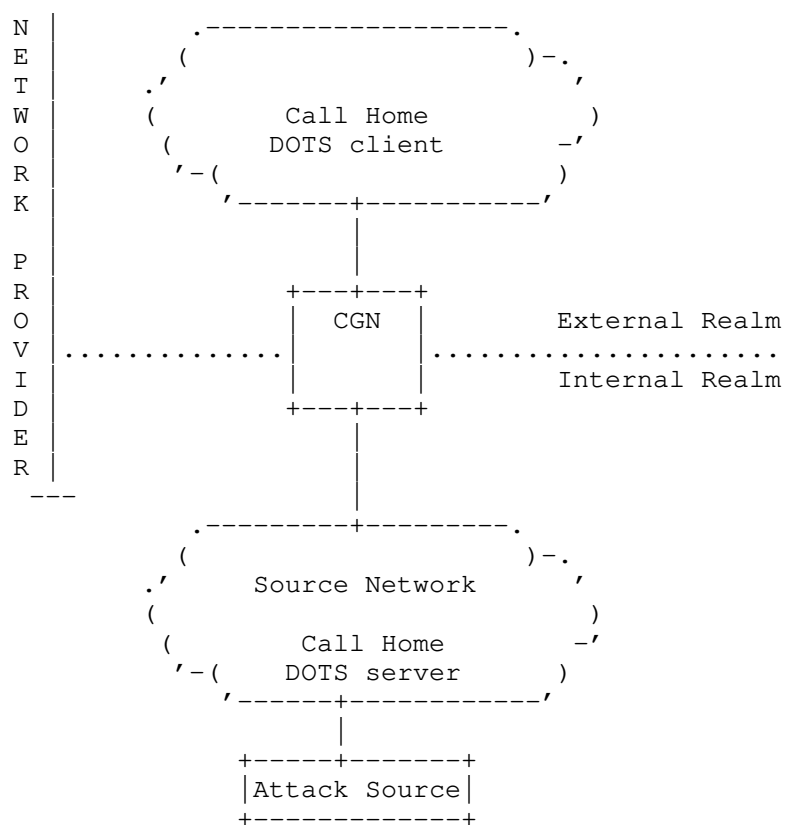


Figure 11: Example of a CGN between DOTS Domains

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 Call Home DOTS server (e.g., a CPE with NAT enabled as shown in Figures 12 and 13), the Call Home 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. The Call Home DOTS server proceeds with a NAT mapping table lookup using the attack information (or a subset thereof) as a key. The lookup can be local (Figure 12) or via a dedicated administration interface that is offered by the CPE

(Figure 13). This identification allows the suspicious device to be isolated while avoiding disturbances of other services.

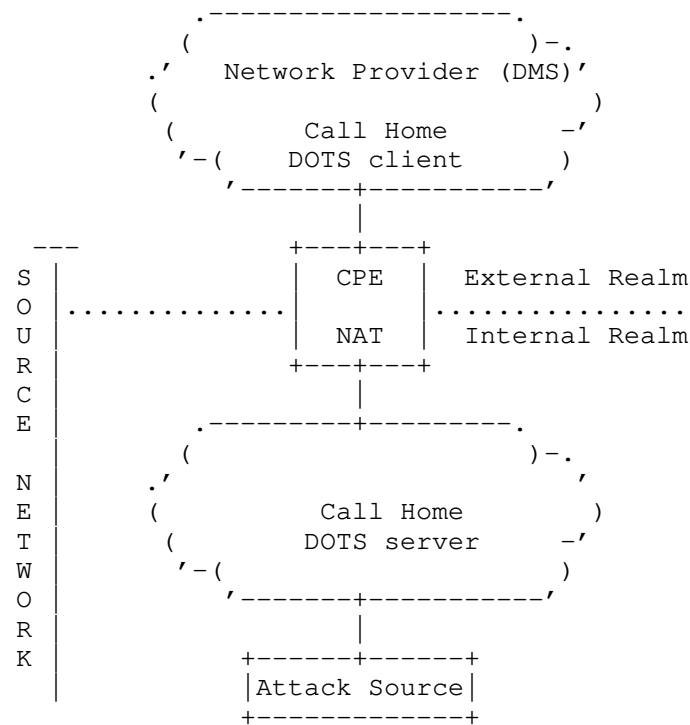


Figure 12: Example of a DOTS Server Domain with a NAT Embedded in a CPE

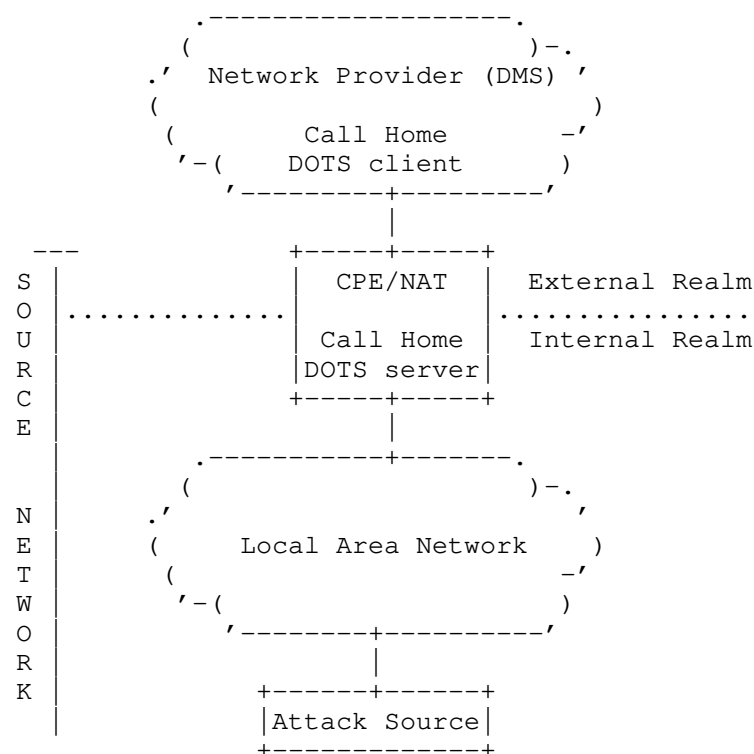


Figure 13: Example of a Call Home DOTS Server and a NAT Embedded in a CPE

If for any reason address sharing is deployed in both source and provider networks, both Call Home DOTS agents have to proceed with address mapping lookups following the behavior specified in reference to Figure 11 (network provider) and Figures 12 and 13 (source network).

6. DOTS Signal Call Home YANG Module

6.1. Tree Structure

This document augments the "ietf-dots-signal-channel" (dots-signal) DOTS signal YANG module defined in [I-D.ietf-dots-rfc8782-bis] 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-structure /dots-signal:dots-signal/dots-signal:message-type
    /dots-signal:mitigation-scope/dots-signal:scope:
  +-- source-prefix*          inet:ip-prefix
  +-- source-port-range* [lower-port]
    | +-- lower-port      inet:port-number
    | +-- upper-port?    inet:port-number
  +-- source-icmp-type-range* [lower-type]
    +-- lower-type      uint8
    +-- upper-type?    uint8
augment-structure /dots-signal:dots-signal/dots-signal:message-type
    /dots-signal:redirected-signal:
  +-- (type)?
    +--:(call-home-only)
      +-- alt-ch-client          inet:domain-name
      +-- alt-ch-client-record*  inet:ip-address
      +-- ttl?                   uint32
```

6.2. YANG/JSON Mapping Parameters to CBOR

The YANG/JSON mapping parameters to CBOR are listed in Table 1.

- o 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	CBOR Major Type & Information	JSON Type
ietf-dots-call-home:source-prefix	leaf-list inet: ip-prefix	TBA1	4 array 3 text string	Array String
ietf-dots-call-home:source-port-range	list	TBA2	4 array	Array
ietf-dots-call-home:source-icmp-type-range	list	TBA3	4 array	Array
lower-type	uint8	TBA4	0 unsigned	Number
upper-type	uint8	TBA5	0 unsigned	Number
ietf-dots-call-home:alt-ch-client	inet: domain-name	TBA6	3 text string	String
ietf-dots-call-home:alt-ch-client-record	leaf-list inet: ip-address	TBA7	4 array 3 text string	Array String
ietf-dots-call-home:ttl	uint32	TBA8	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 [I-D.ietf-dots-rfc8782-bis].

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@2020-12-02.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 YYYY: 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:TirumaleswarReddy\_Konda@McAfee.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 XXXX; see
    the RFC itself for full legal notices.";
```

```
revision 2020-12-02 {
  description
    "Initial revision.";
  reference
    "RFC XXXX: 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 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 are 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 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 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 Internet Protocol Version 6 (IPv6)  

        Specification.";
}
}
}
sx:augment-structure "/dots-signal:dots-signal"  

    + "/dots-signal:message-type"  

    + "/dots-signal:redirected-signal" {
description
    "Augments the redirected 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 call home signal  

        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 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 YYYY";
    }
}
}
}
}
}
<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].

- o Note to the RFC Editor: Please delete TBA1-TBA8 once CBOR keys are assigned from the 32768-49151 range.

Parameter Name	CBOR Key Value	CBOR Major Type	Change Controller	Specification Document (s)
ietf-dots-call-home: source-prefix	TBA1	4	IESG	[RFCXXXX]
ietf-dots-call-home: source-port-range	TBA2	4	IESG	[RFCXXXX]
ietf-dots-call-home: source-icmp-type- range	TBA3	4	IESG	[RFCXXXX]
lower-type	TBA4	0	IESG	[RFCXXXX]
upper-type	TBA5	0	IESG	[RFCXXXX]
ietf-dots-call-home: alt-ch-client	TBA6	3	IESG	[RFCXXXX]
ietf-dots-call-home: alt-ch-client-record	TBA7	4	IESG	[RFCXXXX]
ietf-dots-call-home: ttl	TBA8	0	IESG	[RFCXXXX]

Table 2: Assigned DOTS Signal Channel CBOR Key Values

7.2. New DOTS Conflict Cause

This document requests IANA to assign a new code from the "DOTS Signal Channel Conflict Cause Codes" registry [Cause].

Code	Label	Description	Reference
4 (TB A9)	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.	[RFCXXXX]

7.3. DOTS Signal Call Home YANG Module

This document requests IANA to register 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.

This document requests IANA to register 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 XXXX

8. Security Considerations

This document deviates from classic DOTS signal channel usage by having the DOTS server initiate the (D)TLS connection. DOTS signal channel related security considerations discussed in Section 11 of [I-D.ietf-dots-rfc8782-bis] 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 to a drop-list the source address after a set number of unsuccessful authentication attempts.

The DOTS Call Home signal channel can be misused by a misbehaving Call Home DOTS client by arbitrarily signalling 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, or could re-direct 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 DOTS server domain administrator to block the traffic from the potentially compromised device to the target (see Section 5.3.1). Means to seek the 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]).

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 learnt 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 to (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 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 a given attack detection technology. The reader may refer, for example, to Section 11.8 of [RFC7011].

10. 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

Wei Pan
Huawei Technologies
China

Email: william.panwei@huawei.com

11. Acknowledgements

Thanks to Wei Pei, Xia Liang, Roman Danyliw, Dan Wing, Toema Gavrichenkov, Daniel Migault, 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 yangdoctors review.

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

12. References

12.1. Normative References

- [I-D.ietf-dots-rfc8782-bis]
Boucadair, M., Shallow, J., and T. Reddy.K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-rfc8782-bis-06 (work in progress), March 2021.
- [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., Bjoerklund, M., and K. Watsen, "YANG Data Structure Extensions", RFC 8791, DOI 10.17487/RFC8791, June 2020, <<https://www.rfc-editor.org/info/rfc8791>>.

12.2. Informative References

- [Cause] IANA, "DOTS Signal Channel Conflict Cause Codes", <<https://www.iana.org/assignments/dots/dots.xhtml#dots-signal-channel-conflict-cause-codes>>.
- [I-D.ietf-dots-multihoming] Boucadair, M., Reddy, T., and W. Pan, "Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS)", draft-ietf-dots-multihoming-05 (work in progress), November 2020.
- [I-D.ietf-dots-use-cases] Dobbins, R., Migault, D., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", draft-ietf-dots-use-cases-25 (work in progress), July 2020.
- [I-D.ietf-i2nsf-terminology] Hares, S., Strassner, J., Lopez, D. R., Xia, L., and H. Birkholz, "Interface to Network Security Functions (I2NSF) Terminology", draft-ietf-i2nsf-terminology-08 (work in progress), July 2019.
- [Key-Map] IANA, "DOTS Signal Channel CBOR Key Values", <<https://www.iana.org/assignments/dots/dots.xhtml#dots-signal-channel-cbor-key-values>>.
- [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>>.

- [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://hackers.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 [I-D.ietf-i2nsf-terminology] 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 [I-D.ietf-dots-rfc8782-bis] 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 [I-D.ietf-dots-rfc8782-bis] 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].

Authors' Addresses

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

Email: kondtir@gmail.com

Mohamed Boucadair (editor)
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Jon Shallow
UK

Email: supjps-ietf@jpshallow.com

DOTS
Internet-Draft
Intended status: Informational
Expires: 4 October 2022

Y. Hayashi
NTT
M. Chen
Li. Su
CMCC
2 April 2022

Use Cases for DDoS Open Threat Signaling (DOTS) Telemetry
draft-ietf-dots-telemetry-use-cases-10

Abstract

DDoS Open Threat Signaling (DOTS) Telemetry enriches the base DOTS protocols to assist the mitigator in using efficient DDoS attack mitigation techniques in a network. This document presents sample use cases for DOTS Telemetry. It discusses in particular what components are deployed in the network, how they cooperate, and what information is exchanged to effectively use these techniques.

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 4 October 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Terminology	3
3. Telemetry Use Cases	3
3.1. Mitigation Resources Assignment	3
3.1.1. Mitigating Attack Flow of Top-talker Preferentially	3
3.1.2. Optimal DMS Selection for Mitigation	6
3.1.3. Best-path Selection for Redirection	9
3.1.4. Short but Extreme Volumetric Attack Mitigation	11
3.1.5. Selecting Mitigation Technique Based on Attack Type	14
3.2. Detailed DDoS Mitigation Report	18
3.3. Tuning Mitigation Resources	21
3.3.1. Supervised Machine Learning of Flow Collector	21
3.3.2. Unsupervised Machine Learning of Flow Collector	24
4. Security Considerations	26
5. IANA Considerations	26
6. Acknowledgement	26
7. References	26
7.1. Normative References	26
7.2. Informative References	26
Authors' Addresses	28

1. Introduction

Distributed Denial-of-Service (DDoS) attacks, such as volumetric attacks and resource-consumption attacks, are critical threats to be handled by service providers. When such DDoS attacks occur, service providers have to mitigate them immediately to protect or recover their services.

Therefore, for service providers to immediately protect their network services from DDoS attacks, DDoS mitigation needs to be highly automated. To that aim, multi-vendor components involved in DDoS attack detection and mitigation should cooperate and support standard interfaces.

DDoS Open Threat Signaling (DOTS) is a set of protocols for real-time signaling, threat-handling requests, and data filtering between the multi-vendor elements [RFC9132][RFC8783]. DOTS Telemetry enriches the DOTS protocols with various telemetry attributes allowing optimal DDoS attack mitigation [I-D.ietf-dots-telemetry]. This document presents sample use cases for DOTS Telemetry, which makes concrete overview and purpose described in [I-D.ietf-dots-telemetry]: what components are deployed in the network, how they cooperate, and what information is exchanged to effectively use attack-mitigation techniques.

2. Terminology

The readers should be familiar with the terms defined in [RFC8612], [RFC8903] and [I-D.ietf-dots-telemetry].

In addition, this document uses the following terms:

Top-talker: A list of attack sources that are involved in an attack and which are generating an important part of the attack traffic.

Supervised Machine Learning: A machine-learning technique in which labeled data is used to train the algorithms (the input and output data are known).

Unsupervised Machine Learning: A machine learning technique in which unlabeled data is used to train the algorithms (the data has no historical labels).

3. Telemetry Use Cases

This section describes DOTS telemetry use cases that use attributes included in DOTS telemetry specifications [I-D.ietf-dots-telemetry].

3.1. Mitigation Resources Assignment


3.1.1. Mitigating Attack Flow of Top-talker Preferentially

Some transit providers have to mitigate such large-scale DDoS attacks by using DDoS Mitigation Systems (DMSes) with limited resources, which is already deployed in their network. For example, recent reported large DDoS attacks exceeded 1 Tps.

The aim of this use case is to enable transit providers to use their DMS efficiently under volume-based DDoS attacks whose volume is more than the available capacity of the DMS. To enable this, the attack traffic of top-talkers is redirected to the DMS preferentially by cooperation among forwarding nodes, flow collectors, and orchestrators.

Figure 1 gives an overview of this use case. Figure 2 provides an example of a DOTS telemetry message body that is used to signal top-talkers (2001:db8::2/128 and 2001:db8::3/128).

(Internet Transit Provider)



* C is for DOTS client functionality

* S is for DOTS server functionality

Figure 1: Mitigating DDoS Attack Flow of Top-talkers Preferentially

```
{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "total-attack-traffic-protocol": [
          {
            "protocol": 17,
            "unit": "megabit-ps",
            "mid-percentile-g": "900"
          }
        ],
        "attack-detail": [
          {
            "vendor-id": 32473,
            "attack-id": 77,
            "start-time": "1645057211",
            "attack-severity": "high",
            "top-talker": {
              "talker": [
                {
                  "source-prefix": "2001:db8::2/128",
                  "total-attack-traffic": [
                    {
                      "unit": "megabit-ps",
                      "mid-percentile-g": "100"
                    }
                  ]
                },
                {
                  "source-prefix": "2001:db8::3/128",
                  "total-attack-traffic": [
                    {
                      "unit": "megabit-ps",
                      "mid-percentile-g": "90"
                    }
                  ]
                }
              ]
            }
          }
        ]
      }
    ]
  }
}
```

```
}  
}
```

Figure 2: An Example of Message Body to Signal Top-Talkers

The forwarding nodes send traffic statistics to the flow collectors using, e.g., IP Flow Information Export (IPFIX) [RFC7011]. When DDoS attacks occur, the flow collectors identifies the attack traffic and send information of the top-talkers to the orchestrator using the "target-prefix" and "top-talkers" DOTS telemetry attributes. The orchestrator, then, checks the available capacity of the DMSes by using a network management protocol, such as Simple Network Management Protocol (SNMP) [RFC3413]. After that, the orchestrator orders the forwarding nodes to redirect as much of the top-talker's traffic to the DMS as possible by dissemination of Flow Specifications relying upon tools, such as Border Gateway Protocol Dissemination of Flow Specification Rules (BGP Flowspec) [RFC8955].

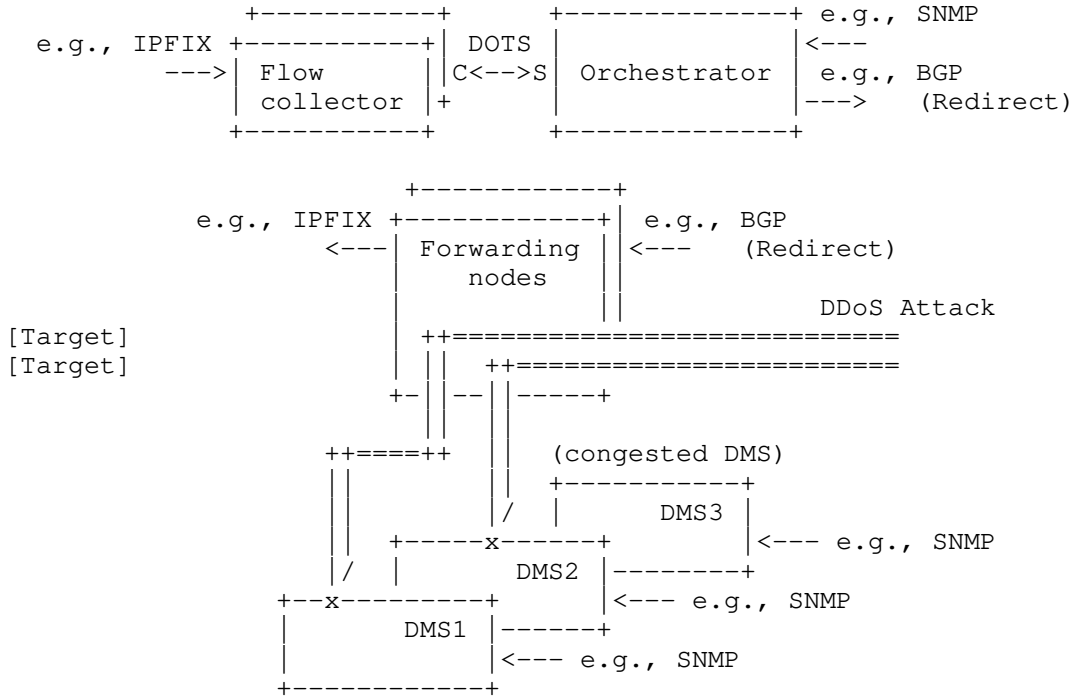
The flow collector implements a DOTS client while the orchestrator implements a DOTS server.

3.1.2. Optimal DMS Selection for Mitigation

Transit providers can deploy their DMSes in clusters. Then, they can select the DMS to be used to mitigate a DDoS attack under attack time.

The aim of this use case is to enable transit providers to select an optimal DMS for mitigation based on the volume of the attack traffic and the capacity of a DMS. Figure 3 gives an overview of this use case. Figure 4 provides an example of a DOTS telemetry message body that is used to signal various attack traffic percentiles.

(Internet Transit Provider)



- * C is for DOTS client functionality
- * S is for DOTS server functionality

Figure 3: Optimal DMS Selection for Mitigation

```
{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "total-attack-traffic": [
          {
            "unit": "megabit-ps",
            "low-percentile-g": "600",
            "mid-percentile-g": "800",
            "high-percentile-g": "1000",
            "peak-g": "1100",
            "current-g": "700"
          }
        ]
      }
    ]
  }
}
```

Figure 4: Example of Message Body with Total Attack Traffic

The forwarding nodes send traffic statistics to the flow collectors using, e.g., IPFIX. When DDoS attacks occur, the flow collectors identify the attack traffic and send information of the attack traffic volume to the orchestrator by using the "target-prefix" and "total-attack-traffic" DOTS telemetry attributes. The orchestrator, then, checks the available capacity of the DMSes by using a network management protocol, such as SNMP. After that, the orchestrator selects an optimal DMS to which each attack traffic should be redirected. For example, a simple DMS selection algorithm is to choose a DMS whose available capacity is greater than the "peak-g" attribute indicated in the DOTS telemetry message. The orchestrator orders the appropriate forwarding nodes to redirect the attack traffic to the optimal DMS relying upon routing policies, such as BGP [RFC4271].

The detailed DMS selection algorithm is out of the scope of this document.

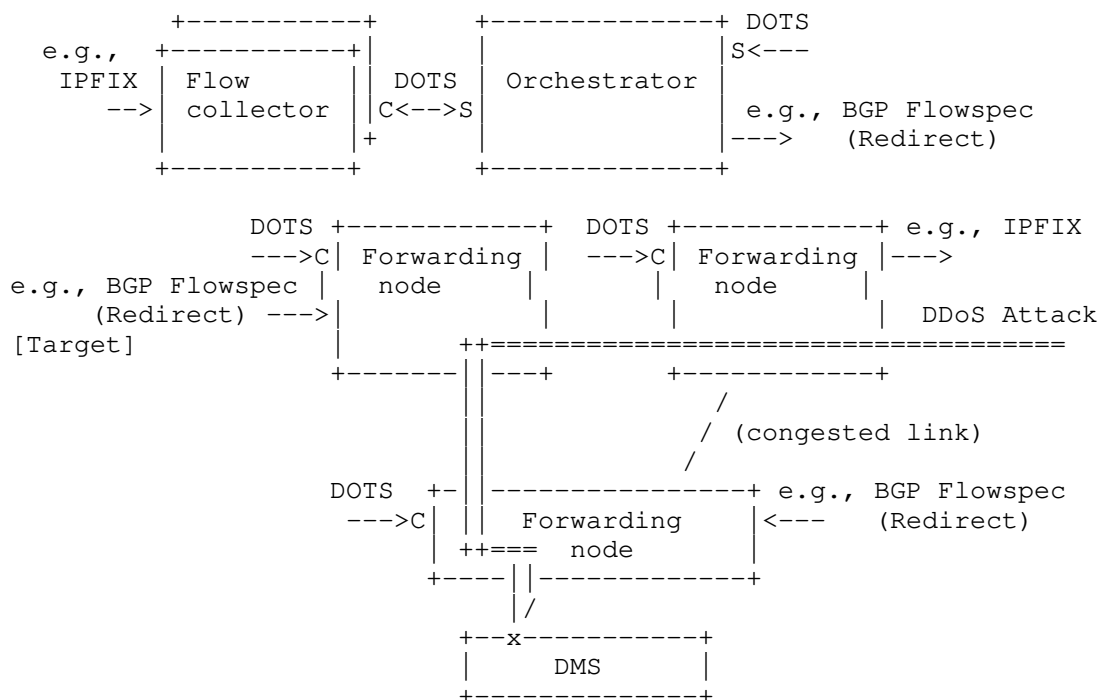
The flow collector implements a DOTS client while the orchestrator implements a DOTS server.

3.1.3. Best-path Selection for Redirection

A transit provider network has multiple paths to convey an attack traffic to a DMS. In such a network, the attack traffic can be conveyed while avoiding congested links by adequately selecting an available path.

The aim of this use case is to enable transit providers to select an optimal path for redirecting attack traffic to a DMS according to the bandwidth of the attack traffic and total traffic. Figure 5 gives an overview of this use case. Figure 6 provides an example of a DOTS telemetry message body that is used to signal various attack traffic percentiles and total traffic percentiles.

(Internet Transit Provider)



* C is for DOTS client functionality

* S is for DOTS server functionality

Figure 5: Best-path Selection for Redirection


```
{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "total-traffic": [
          {
            "unit": "megabit-ps",
            "mid-percentile-g": "1300",
            "peak-g": "800"
          }
        ],
        "total-attack-traffic": [
          {
            "unit": "megabit-ps",
            "low-percentile-g": "600",
            "mid-percentile-g": "800",
            "high-percentile-g": "1000",
            "peak-g": "1100",
            "current-g": "700"
          }
        ]
      }
    ]
  }
}
```

Figure 6: An Example of Message Body with Total Attack Traffic and Total Traffic

The forwarding nodes send traffic statistics to the flow collectors by using, e.g., IPFIX. When DDoS attacks occur, the flow collectors identify attack traffic and send information of the attack traffic volume to the orchestrator by using a "target-prefix" and "total-attack-traffic" DOTS telemetry attributes. On the other hands, the underlying forwarding nodes send volume of the total traffic passing the node to the orchestrator by using "total-traffic" telemetry attributes. The orchestrator then selects an optimal path to which each attack-traffic flow should be redirected. For example, the simple algorithm of the selection is to choose a path whose available capacity is greater than the "peak-g" attribute that was indicated in a DOTS telemetry message. After that, the orchestrator orders the appropriate forwarding nodes to redirect the attack traffic to the optimal DMS by dissemination of Flow Specifications relying upon tools, such as BGP Flowspec.

The detailed path selection algorithm is out of the scope of this document.

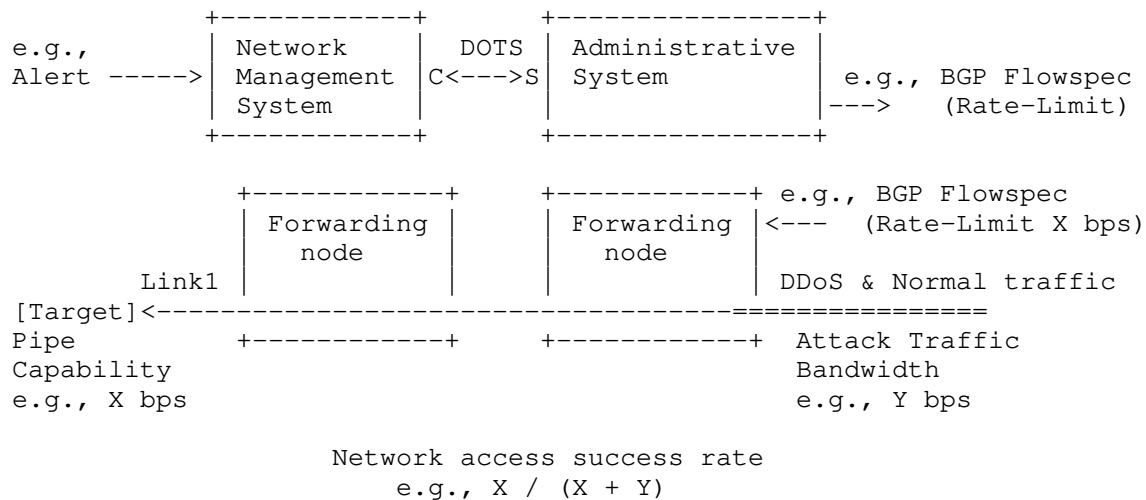
The flow collector and forwarding nodes implement a DOTS client while the orchestrator implements a DOTS server.

3.1.4. Short but Extreme Volumetric Attack Mitigation

Short, but extreme volumetric attacks, such as pulse wave DDoS attacks, are threats to internet transit provider networks. The feature of the attack is that start from zero and go to maximum values in a very short time span, then go back to zero, and back to maximum, repeating in continuous cycles at short intervals. It is difficult for them to mitigate an attack by DMS by redirecting attack flows because it may cause route flapping in the network. The practical way to mitigate short but extreme volumetric attacks is to offload mitigation actions to a forwarding node.

The aim of this use case is to enable transit providers to mitigate short but extreme volumetric attacks. Furthermore, the aim is to estimate the network-access success rate based on the bandwidth of attack traffic. Figure 7 gives an overview of this use case. Figure 8 provides an example of a DOTS telemetry message body that is used to signal total pipe capacity. Figure 9 provides an example of a DOTS telemetry message body that is used to signal various attack traffic percentiles and total traffic percentiles.

(Internet Transit Provider)



* C is for DOTS client functionality

* S is for DOTS server functionality

Figure 7: Short but Extreme Volumetric Attack Mitigation

```
{
  "ietf-dots-telemetry:telemetry-setup": {
    "telemetry": [
      {
        "total-pipe-capacity": [
          {
            "link-id": "link1",
            "capacity": "1000",
            "unit": "megabit-ps"
          }
        ]
      }
    ]
  }
}
```

Figure 8: Example of Message Body with Total Pipe Capacity

```
{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "total-traffic": [
          {
            "unit": "megabit-ps",
            "mid-percentile-g": "800",
            "peak-g": "1300"
          }
        ],
        "total-attack-traffic": [
          {
            "unit": "megabit-ps",
            "low-percentile-g": "200",
            "mid-percentile-g": "400",
            "high-percentile-g": "500",
            "peak-g": "600",
            "current-g": "400"
          }
        ]
      }
    ]
  }
}
```

Figure 9: Example of Message Body with Total Attack Traffic,
and Total Traffic

When DDoS attacks occur, the network management system receives alerts. Then, it sends the target IP address(es) and volume of the DDoS attack traffic to the administrative system by using the "target-prefix" and "total-attack-traffic" DOTS telemetry attributes. After that, the administrative system orders relevant forwarding nodes to carry out rate-limit all traffic destined to the target based on the pipe capability by the dissemination of the Flow Specifications relying upon tools, such as BGP Flowspec. In addition, the administrative system estimates the network-access success rate of the target, which is calculated by $(\text{total-pipe-capability} / (\text{total-pipe-capability} + \text{total-attack-traffic}))$.

Note that total pipe capability information can be gathered by telemetry setup in advance (Section 7.2 of [I-D.ietf-dots-telemetry]).

The network management system implements a DOTS client while the administrative system implements a DOTS server.

3.1.5. Selecting Mitigation Technique Based on Attack Type

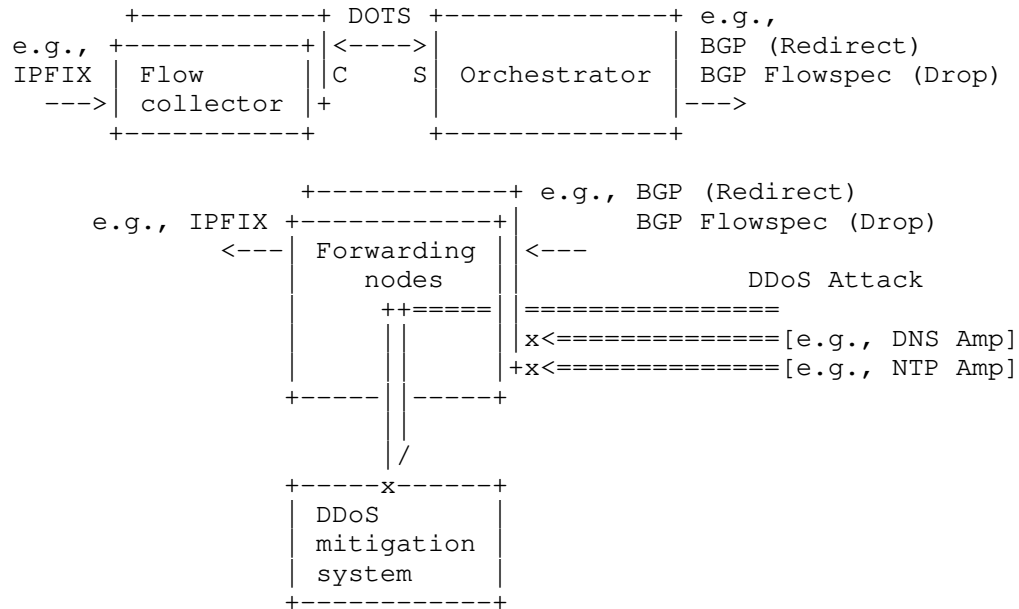
Some volumetric attacks, such as amplification attacks, can be detected with high accuracy by checking the Layer 3 or Layer 4 information of attack packets. These attacks can be detected and mitigated through cooperation among forwarding nodes and flow collectors by using IPFIX. It may also be necessary to inspect the Layer 7 information of suspicious packets to detect attacks such as DNS Water Torture Attacks. Such an attack traffic should be detected and mitigated at a DMS.

The aim of this use case is to enable transit providers to select a mitigation technique based on the type of attack traffic: amplification attack or not. To use such a technique, the attack traffic is blocked by forwarding nodes or redirected to a DMS based on the attack type through cooperation among forwarding nodes, flow collectors, and an orchestrator.

Figure 10 gives an overview of this use case. Figure 11 provides an example of attack mappings as below are shared by using the DOTS data channel in advance. Figure 12 provides an example of a DOTS telemetry message body that is used to signal various attack traffic percentiles, total traffic percentiles, total attack connection and attack type.

The example in Figure 11 uses the folding defined in [RFC8792] for long lines.

(Internet Transit Provider)



- * C is for DOTS client functionality
- * S is for DOTS server functionality
- * DNS Amp: DNS Amplification
- * NTP Amp: NTP Amplification

Figure 10: DDoS Mitigation Based on Attack Type

===== NOTE: '\ ' line wrapping per RFC 8792 =====

```
{
  "ietf-dots-mapping:vendor-mapping": {
    "vendor": [
      {
        "vendor-id": 32473,
        "vendor-name": "mitigator-c",
        "last-updated": "1629898958",
        "attack-mapping": [
          {
            "attack-id": 77,
            "attack-description":
              "attack-description": "DNS amplification Attack: \
This attack is a type of reflection attack in which attackers \
spooof a target's IP address. The attackers abuse vulnerabilities \
in DNS servers to turn small queries into larger payloads."
          },
          {
            "attack-id": 92,
            "attack-description":
              "attack-description": "NTP amplification Attack: \
This attack is a type of reflection attack in which attackers \
spooof a target's IP address. The attackers abuse vulnerabilities \
in NTP servers to turn small queries into larger payloads."
          }
        ]
      }
    ]
  }
}
```

Figure 11: Example of Message Body with Attack Mappings

```
{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "total-attack-traffic": [
          {
            "unit": "megabit-ps",
            "low-percentile-g": "600",
            "mid-percentile-g": "800",

```

```
        "high-percentile-g": "1000",
        "peak-g": "1100",
        "current-g": "700"
      }
    ],
    "total-attack-traffic-protocol": [
      {
        "protocol": 17,
        "unit": "megabit-ps",
        "mid-percentile-g": "500"
      },
      {
        "protocol": 15,
        "unit": "megabit-ps",
        "mid-percentile-g": "200"
      }
    ],
    "total-attack-connection": [
      {
        "mid-percentile-l": [
          {
            "protocol": 15,
            "connection": 200
          }
        ],
        "high-percentile-l": [
          {
            "protocol": 17,
            "connection": 300
          }
        ]
      }
    ],
    "attack-detail": [
      {
        "vendor-id": 32473,
        "attack-id": 77,
        "start-time": "1641169211",
        "attack-severity": "high"
      },
      {
        "vendor-id": 32473,
        "attack-id": 92,
        "start-time": "1641172809",
        "attack-severity": "high"
      }
    ]
  }
}
```



```
    ]  
  }  
}
```

Figure 12: Example of Message Body with Total Attack Traffic, Total Attack Traffic Protocol, Total Attack Connection and Attack Type

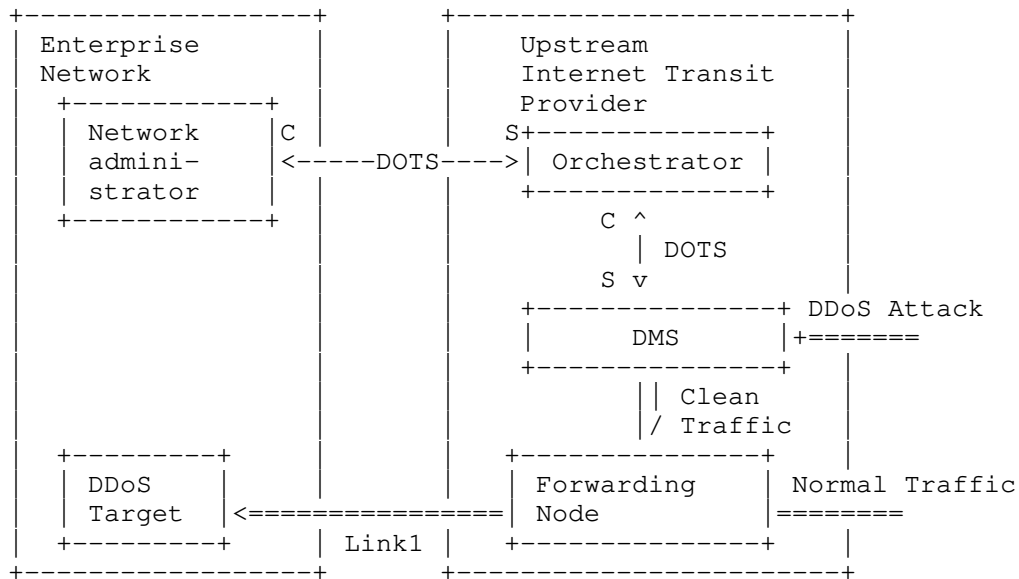
Attack mappings are shared by using the DOTS data channel in advance (Section 8.1.6 of [I-D.ietf-dots-telemetry]). The forwarding nodes send traffic statistics to the flow collectors by using, e.g., IPFIX. When DDoS attacks occur, the flow collectors identify attack traffic and send attack type information to the orchestrator by using "vendor-id" and "attack-id" telemetry attributes. The orchestrator, then, resolves abused port numbers and orders relevant forwarding nodes to block the amplification attack traffic flow by dissemination of Flow Specifications, e.g. [RFC8955]. Also, the orchestrator orders relevant forwarding nodes to redirect other traffic than the amplification attack traffic by using a routing protocol, such as BGP.

The flow collector implements a DOTS client while the orchestrator implements a DOTS server.

3.2. Detailed DDoS Mitigation Report

It is possible for the transit provider to add value to the DDoS mitigation service by reporting on-going and detailed DDoS countermeasure status to the enterprise network. In addition, it is possible for the transit provider to know whether the DDoS counter measure is effective or not by receiving reports from the enterprise network.

The aim of this use case is to share the information about on-going DDoS counter measure between the transit provider and the enterprise network mutually. Figure 13 gives an overview of this use case. Figure 14 provides an example of a DOTS telemetry message body that is used to signal total pipe capacity from the enterprise network administrator to the orchestrator in the ISP. Figure 15 provides an example of a DOTS telemetry message body that is used to signal various total traffic percentiles, total attack traffic percentiles and attack detail from the orchestrator to the network.



* C is for DOTS client functionality
 * S is for DOTS server functionality

Figure 13: Detailed DDoS Mitigation Report

```

{
  "ietf-dots-telemetry:telemetry-setup": {
    "telemetry": [
      {
        "total-pipe-capacity": [
          {
            "link-id": "link1",
            "capacity": "1000",
            "unit": "megabit-ps"
          }
        ]
      }
    ]
  }
}

```

Figure 14: An Example of Message Body with Total Pipe Capacity

```

{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "tmid": 567,
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "target-protocol": [
          17
        ],
        "total-traffic": [
          {
            "unit": "megabit-ps",
            "mid-percentile-g": "800"
          }
        ],
        "total-attack-traffic": [
          {
            "unit": "megabit-ps",
            "mid-percentile-g": "100"
          }
        ],
        "attack-detail": [
          {
            "vendor-id": 32473,
            "attack-id": 77,
            "start-time": "1644819611",
            "attack-severity": "high"
          }
        ]
      }
    ]
  }
}

```

Figure 15: An Example of Message Body with Total Traffic,
Total Attack Traffic Protocol, and Attack Detail

The network management system in the enterprise network reports limits of incoming traffic volume from the transit provider to the orchestrator in the transit provider in advance. It is reported by using "total-pipe-capacity" telemetry attribute in DOTS telemetry setup.

When DDoS attacks occur, DDoS mitigation orchestration [RFC8903] is carried out in the transit provider. Then, the DDoS mitigation systems reports the status of DDoS countermeasures to the orchestrator by sending "attack-detail" telemetry attributes. After that, the orchestrator integrates the reports from the DDoS mitigation system, while removing duplicate contents, and sends them to a network administrator by using DOTS telemetry periodically.

During the DDoS mitigation, the orchestrator in the transit provider retrieves link congestion status from the network manager in the enterprise network by using "total-traffic" telemetry attributes. Then, the orchestrator checks whether the DDoS countermeasures are effective or not by comparing the "total-traffic" and the "total-pipe-capacity" attributes.

The DMS implements a DOTS server while the orchestrator behaves as a DOTS client and a server in the transit provider. In addition, the network administrator implements a DOTS client.

3.3. Tuning Mitigation Resources

3.3.1. Supervised Machine Learning of Flow Collector

DDoS detection based on tools, such as IPFIX, is a lighter weight method of detecting DDoS attacks than DMSes in internet transit provider networks. On the other hand, DDoS detection based on the DMSes is a more accurate method for detecting attack traffic than flow monitoring.

The aim of this use case is to increase flow collector's detection accuracy by carrying out supervised machine-learning techniques according to attack detail reported by the DMSes. To use such a technique, forwarding nodes, flow collector, and a DMS should cooperate. Figure 16 gives an overview of this use case. Figure 17 provides an example of a DOTS telemetry message body that is used to signal various total attack traffic percentiles and attack detail.

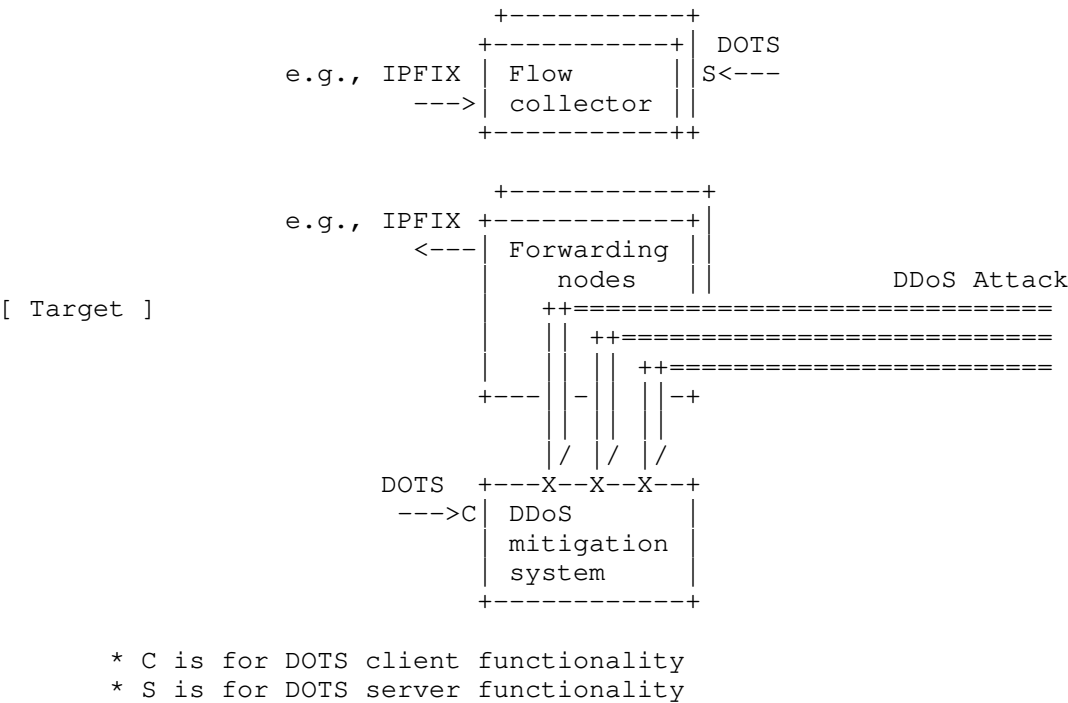


Figure 16: Training Supervised Machine Learning of Flow Collectors

```

{
  "ietf-dots-telemetry:telemetry": {
    "pre-or-ongoing-mitigation": [
      {
        "target": {
          "target-prefix": [
            "2001:db8::1/128"
          ]
        },
        "attack-detail": [
          {
            "vendor-id": 32473,
            "attack-id": 77,
            "start-time": "1634192411",
            "attack-severity": "high",
            "top-talker": {
              "talker": [
                {
                  "source-prefix": "2001:db8::2/128"
                },
                {
                  "source-prefix": "2001:db8::3/128"
                }
              ]
            }
          }
        ]
      }
    ]
  }
}

```

Figure 17: An Example of Message Body with Attack Type and top-talkers

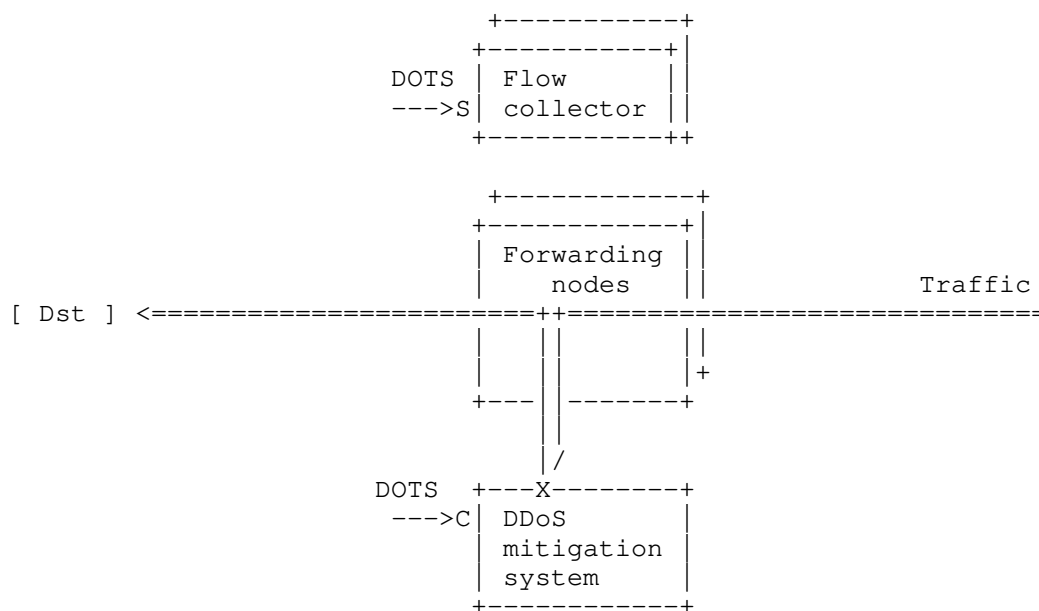
The forwarding nodes send traffic statistics to the flow collectors by using, e.g., IPFIX. When DDoS attacks occur, DDoS mitigation orchestration is carried out (as per Section 3.3 of [RFC8903]) and the DMS mitigates all attack traffic destined for a target. The DDoS mitigation system reports the "vendor-id", "attack-id", and "top-talker" telemetry attributes to a flow collector.

After mitigating a DDoS attack, the flow collector attaches outputs of the DMS as labels to the statistics of traffic flow of top-talkers. The outputs, for example, are the "attack-id" telemetry attributes. The flow collector, then, carries out supervised machine learning to increase its detection accuracy, setting the statistics as an explanatory variable and setting the labels as an objective variable.

The DMS implements a DOTS client while the flow collector implements a DOTS server.

3.3.2. Unsupervised Machine Learning of Flow Collector

DMSes can detect DDoS attack traffic, which means DMSes can also identify clean traffic. The aim of this use case is to carry out unsupervised machine-learning for anomaly detection according to baseline reported by DMSes. To use such a technique, forwarding nodes, flow collector, and a DMS should cooperate. Figure 18 gives an overview of this use case. Figure 19 provides an example of a DOTS telemetry message body that is used to signal baseline.



- * C is for DOTS client functionality
- * S is for DOTS server functionality

Figure 18: Training Unsupervised Machine Learning of Flow Collectors

```

{
  "ietf-dots-telemetry:telemetry-setup": {
    "telemetry": [
      {
        "baseline": [
          {
            "id": 1,
            "target-prefix": [
              "2001:db8:6401::1/128"
            ],
            "target-port-range": [
              {
                "lower-port": "53"
              }
            ],
            "target-protocol": [
              17
            ],
            "total-traffic-normal": [
              {
                "unit": "megabit-ps",
                "mid-percentile-g": "30",
                "mid-percentile-g": "50",
                "high-percentile-g": "60",
                "peak-g": "70"
              }
            ]
          }
        ]
      }
    ]
  }
}

```

Figure 19: An Example of Message Body with Traffic Baseline

The forwarding nodes carry out mirroring traffic destined IP address. The DMS then identifies "clean" traffic and reports the baseline attributes to the flow collector by using DOTS telemetry.

The flow collector, then, carries out unsupervised machine learning to be able to carry out anomaly detection.

The DMS implements a DOTS client while the flow collector implements a DOTS server.

4. Security Considerations

DOTS telemetry security considerations are discussed in Section 14 of [I-D.ietf-dots-telemetry]. These considerations apply for the communication interfaces where DOTS is used.

Some use cases involve controllers, orchestrators, and programmable interfaces. These interfaces can be misused by misbehaving nodes to further exacerbate DDoS attacks. Section 5 of [RFC7149] discusses some generic security considerations to take into account in such contexts (e.g., reliable access control). Specific security measures depend on the actual mechanism used to control underlying forwarding nodes and other controlled elements. For example, Section 13 of [RFC8955] discusses security considerations that are relevant to BGP Flowspec. IPFIX-specific considerations are discussed in Section 11 of [RFC7011].

5. IANA Considerations

This document does not require any action from IANA.

6. Acknowledgement

The authors would like to thank Mohamed Boucadair and Valery Smyslov for their valuable feedback.

7. References

7.1. Normative References

[I-D.ietf-dots-telemetry]
Boucadair, M., Reddy, K. T., Doron, E., Chen, M., and J. Shallow, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Telemetry", Work in Progress, Internet-Draft, draft-ietf-dots-telemetry-25, 21 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-dots-telemetry-25.txt>>.

7.2. Informative References

[RFC3413] Levi, D., Meyer, P., and B. Stewart, "Simple Network Management Protocol (SNMP) Applications", STD 62, RFC 3413, DOI 10.17487/RFC3413, December 2002, <<https://www.rfc-editor.org/info/rfc3413>>.

- [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>>.
- [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>>.
- [RFC7149] Boucadair, M. and C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment", RFC 7149, DOI 10.17487/RFC7149, March 2014, <<https://www.rfc-editor.org/info/rfc7149>>.
- [RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <<https://www.rfc-editor.org/info/rfc8612>>.
- [RFC8783] Boucadair, M., Ed. and T. Reddy.K, Ed., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", RFC 8783, DOI 10.17487/RFC8783, May 2020, <<https://www.rfc-editor.org/info/rfc8783>>.
- [RFC8792] Watsen, K., Auerswald, E., Farrel, A., and Q. Wu, "Handling Long Lines in Content of Internet-Drafts and RFCs", RFC 8792, DOI 10.17487/RFC8792, June 2020, <<https://www.rfc-editor.org/info/rfc8792>>.
- [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>>.
- [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>>.

Authors' Addresses

Yuhei Hayashi
NTT
3-9-11, Midori-cho, Tokyo
180-8585
Japan
Email: yuuhei.hayashi@gmail.com

Meiling Chen
CMCC
32, Xuanwumen West
BeiJing
BeiJing, 100053
China
Email: chenmeiling@chinamobile.com

Li Su
CMCC
32, Xuanwumen West
BeiJing, BeiJing
100053
China
Email: sul@chinamobile.com