

DOTS
Internet-Draft
Intended status: Standards Track
Expires: January 6, 2020

T. Reddy
McAfee
M. Boucadair
Orange
E. Doron
Radware Ltd.
July 5, 2019

Distributed Denial-of-Service Open Threat Signaling (DOTS) Telemetry
draft-reddy-dots-telemetry-00

Abstract

This document aims to enrich DOTS signal channel protocol with various telemetry attributes allowing optimal DDoS attack mitigation. This document specifies the normal traffic baseline and attack traffic telemetry attributes a DOTS client can convey to its DOTS server in the mitigation request, the mitigation status telemetry attributes a DOTS server can communicate to a DOTS client, and the mitigation efficacy telemetry attributes a DOTS client can communicate to a DOTS server. The telemetry attributes can assist the mitigator to choose the DDoS mitigation techniques and perform optimal DDoS attack mitigation.

Status of This Memo

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

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

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

This Internet-Draft will expire on January 6, 2020.

Copyright Notice

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

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

Table of Contents

1. Introduction	2
2. Terminology	4
3. DOTS Telemetry: Overview & Purpose	5
4. DOTS Telemetry Attributes	8
4.1. Pre-mitigation DOTS Telemetry Attributes	8
4.1.1. Total Traffic Normal Baseline	8
4.1.2. Total Pipe Capability	9
4.1.3. Total Attack Traffic	9
4.1.4. Total Traffic	9
4.1.5. Attack Details	9
4.2. DOTS Client to Server Mitigation Efficacy DOTS Telemetry Attributes	10
4.2.1. Total Attack Traffic	10
4.2.2. Attack Details	10
4.3. DOTS Server to Client Mitigation Status DOTS Telemetry Attributes	10
4.3.1. Mitigation Status	10
5. DOTS Telemetry YANG Module	10
5.1. Tree Structure	10
5.2. YANG Module	11
6. IANA Considerations	11
6.1. DOTS Signal Channel CBOR Mappings Registry	11
6.2. DOTS Signal Telemetry YANG Module	11
7. Security Considerations	12
8. Acknowledgements	12
9. References	12
9.1. Normative References	12
9.2. Informative References	13
Authors' Addresses	13

1. Introduction

The Internet security 'battle' between the adversary and security countermeasures is an everlasting one. DDoS attacks have become more vicious and sophisticated in almost all aspects of their maneuvers and malevolent intentions. IT organizations and service providers

are facing DDoS attacks that fall into two broad categories: Network/Transport layer attacks and Application layer attacks. Network/Transport layer attacks target the victim's infrastructure. These attacks are not necessarily aimed at taking down the actual delivered services, but rather to eliminate various network elements (routers, switches, firewalls, transit links, and so on) from serving legitimate user traffic. The main method of such attacks is to send a large volume or high PPS of traffic toward the victim's infrastructure. Typically, attack volumes may vary from a few 100 Mbps/PPS to 100s of Gbps or even Tbps. Attacks are commonly carried out leveraging botnets and attack reflectors for amplification attacks, such as NTP, DNS, SNMP, SSDP, and so on. Application layer attacks target various applications. Typical examples include attacks against HTTP/HTTPS, DNS, SIP, SMTP, and so on. However, all valid applications with their port numbers open at network edges can be attractive attack targets. Application layer attacks are considered more complex and hard to categorize, therefore harder to detect and mitigate efficiently.

To compound the problem, attackers also leverage multi-vector attacks. These merciless attacks are assembled from dynamic attack vectors (Network/Application) and tactics. As such, multiple attack vectors formed by multiple attack types and volumes are launched simultaneously towards a victim. Multi-vector attacks are harder to detect and defend. Multiple and simultaneous mitigation techniques are needed to defeat such attack campaigns. It is also common for attackers to change attack vectors right after a successful mitigation, burdening their opponents with changing their defense methods.

The ultimate conclusion derived from these real scenarios is that modern attacks detection and mitigation are most certainly complicated and highly convoluted tasks. They demand a comprehensive knowledge of the attack attributes, the targeted normal behavior/traffic patterns, as well as the attacker's on-going and past actions. Even more challenging, retrieving all the analytics needed for detecting these attacks is not simple to obtain with the industry's current capabilities.

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

Typically, DOTS clients can be integrated within a DDoS attack detector, or network and security elements that have been actively engaged with ongoing attacks. The DOTS client mitigation environment determines that it is no longer possible or practical for it to handle these attacks. This can be due to lack of resources or security capabilities, as derived from the complexities and the intensity of these attacks. In this circumstance, the DOTS client has invaluable knowledge about the actual attacks that need to be handled by the DOTS server. By enabling the DOTS client to share this comprehensive knowledge of an ongoing attack under specific circumstances, the DOTS server can drastically increase its abilities to accomplish successful mitigation. While the attack is being handled by the DOTS server associated mitigation resources, the DOTS server has the knowledge about the ongoing attack mitigation. The DOTS server can share this information with the DOTS client so that the client can better assess and evaluate the actual mitigation realized.

In some deployments, DOTS clients can send mitigation hints derived from attack details to DOTS servers, with the full understanding that the DOTS server may ignore mitigation hints, as described in [RFC8612] (Gen-004). Mitigation hints will be transmitted across the DOTS signal channel, as the data channel may not be functional during an attack. How a DOTS server is handling normal and attack traffic attributes, and mitigation hints is implementation-specific.

Both DOTS client and server can benefit this information by presenting various information in relevant management, reporting, and portal systems.

This document defines DOTS telemetry attributes the DOTS client can convey to the DOTS server, and vice versa. The DOTS telemetry attributes are not mandatory fields. Nevertheless, when DOTS telemetry attributes are available to a DOTS agent, and absent any policy, it can signal the attributes in order to optimize the overall mitigation service provisioned using DOTS. Some of the DOTS telemetry data are not shared during an attack time.

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 [RFC8612].

"DOTS Telemetry" is defined as the collection of attributes that are used to characterize normal traffic baseline, attacks and their mitigation measures, and any related information that may help in enforcing countermeasures. The DOTS Telemetry is an optional set of attributes that can be signaled in the DOTS signal channel protocol.

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

3. DOTS Telemetry: Overview & Purpose

When signaling a mitigation request, it is most certainly beneficial for the DOTS client to signal to the DOTS server any knowledge regarding ongoing attacks. This can happen in cases where DOTS clients are asking the DOTS server for support in defending against attacks that they have already detected and/or mitigated. These actions taken by DOTS clients are referred to as "signaling the DOTS Telemetry".

If attacks are already detected and categorized by the DOTS client domain, the DOTS server, and its associated mitigation services, can proactively benefit this information and optimize the overall service delivered. It is important to note that DOTS client and server detection and mitigation approaches can be different, and can potentially outcome different results and attack classifications. The DDoS mitigation service treats the ongoing attack details from the client as hints and cannot completely rely or trust the attack details conveyed by the DOTS client.

A basic requirement of security operation teams is to be aware and get visibility into the attacks they need to handle. The DOTS server security operation teams benefit from the DOTS telemetry, especially from the reports of ongoing attacks. Even if some mitigation can be automated, operational teams can use the DOTS telemetry to be prepared for attack mitigation and to assign the correct resources (operation staff, networking and mitigation) for the specific service. Similarly, security operation personnel at the DOTS client side ask for feedback about their requests for protection. Therefore, it is valuable for the DOTS server to share DOTS telemetry with the DOTS client. Thus mutual sharing of information is crucial for "closing the mitigation loop" between the DOTS client and server. For the server side team, it is important to realize that the same attacks that the DOTS server's mitigation resources are seeing are those that the DOTS client is asking to mitigate. For the DOTS client side team, it is important to realize that the DOTS clients receive the required service. For example: understanding that "I asked for mitigation of two attacks and my DOTS server detects and mitigates only one...". Cases of inconsistency in attack

classification between DOTS client and server can be high-lighted, and maybe handled, using the DOTS telemetry attributes.

In addition, management and orchestration systems, at both DOTS client and server sides, can potentially use DOTS telemetry as a feedback to automate various control and management activities derived from ongoing information signaled.

If the DOTS server's mitigation resources have the capabilities to facilitate the DOTS telemetry, the DOTS server adopts its protection strategy and activates the required countermeasures immediately (automation enabled). The overall results of this adoption are optimized attack mitigation decisions and actions.

The DOTS telemetry can also be used to tune the DDoS mitigators with the correct state of the attack. During the last few years, DDoS attack detection technologies have evolved from threshold-based detection (that is, cases when all or specific parts of traffic cross a pre-defined threshold for a certain period of time is considered as an attack) to an "anomaly detection" approach. In anomaly detection, the main idea is to maintain rigorous learning of "normal" behavior and where an "anomaly" (or an attack) is identified and categorized based on the knowledge about the normal behavior and a deviation from this normal behavior. Machine learning approaches are used such that the actual "traffic thresholds" are "automatically calculated" by learning the protected entity normal traffic behavior during peace time. The normal traffic characterization learned is referred to as the "normal traffic baseline". An attack is detected when the victim's actual traffic is deviating from this normal baseline.

In addition, subsequent activities toward mitigating an attack are much more challenging. The ability to distinguish legitimate traffic from attacker traffic on a per packet basis is complex. This complexity originates from the fact that the packet itself may look "legitimate" and no attack signature can be identified. The anomaly can be identified only after detailed statistical analysis. DDoS attack mitigators use the normal baseline during the mitigation of an attack to identify and categorize the expected appearance of a specific traffic pattern. Particularly the mitigators use the normal baseline to recognize the "level of normality" needs to be achieved during the various mitigation process.

Normal baseline calculation is performed based on continuous learning of the normal behavior of the protected entities. The minimum learning period varies from hours to days and even weeks, depending on the protected application behavior. The baseline cannot be learned during active attacks because attack conditions do not characterize the protected entities' normal behavior.

If the DOTS client has calculated the normal baseline of its protected entities, signaling this attribute to the DOTS server along with the attack traffic levels is significantly valuable. The DOTS server benefits from this telemetry by tuning its mitigation resources with the DOTS client's normal baseline. The DOTS server mitigators use the baseline to familiarize themselves with the attack victim's normal behavior and target the baseline as the level of normality they need to achieve. Consequently, the overall mitigation performances obtained are dramatically improved in terms of time to mitigate, accuracy, false-negative, false-positive, and other measures.

Mitigation of attacks without having certain knowledge of normal traffic can be inaccurate at best. This is especially true for recursive signaling (see Section 3.2.3 in [I-D.ietf-dots-use-cases]). In addition, the highly diverse types of use-cases where DOTS clients are integrated also emphasize the need for knowledge of client behavior. Consequently, common global thresholds for attack detection practically cannot be realized. Each DOTS client can have its own levels of traffic and normal behavior. Without facilitating normal baseline signaling, it may be very difficult for DOTS servers in some cases to detect and mitigate the attacks accurately. It is important to emphasize that it is practically impossible for the server's mitigators to calculate the normal baseline, in cases they do not have any knowledge of the traffic beforehand. In addition, baseline learning requires a period of time that cannot be afforded during active attack. Of course, this information can be provided using out-of-band mechanisms or manual configuration at the risk to maintain inaccurate information as the network evolves and "normal" patterns change. The use of a dynamic and collaborative means between the DOTS client and server to identify and share key parameters for the sake of efficient DDoS protect is valuable.

During a high volume attack, DOTS client pipes can be totally saturated. The DOTS client asks the DOTS server to handle the attack upstream so that DOTS client pipes return to a reasonable load level (normal pattern, ideally). At this point, it is essential to ensure that the DOTS server does not overwhelm the DOTS client pipes by sending back "clean traffic", or what it believes is "clean". This can happen when the server has not managed to detect and mitigate all the attacks launched towards the client. In this case, it can be valuable to clients to signal to server the "Total pipe capacity", which is the level of traffic the clients can absorb from the upstream server. Dynamic updating of the condition of pipes between DOTS agents while they are under a DDoS attack is essential. For example, for cases of multiple DOTS clients share the same physical connectivity pipes. It is important to note, that the term "pipe" noted here does not necessarily represent physical pipe, but rather

represents the current level of traffic client can observe from server. The server should activate other mechanisms to ensure it does not saturate the client's pipes unintentionally. The rate-limit action defined in [I-D.ietf-dots-data-channel] can be a reasonable candidate to achieve this objective; the client can ask for the type of traffic (such as ICMP, UDP, TCP port 80) it prefers to limit.

To summarize, timely and effective signaling of up-to-date DOTS telemetry to all elements involved in the mitigation process is essential and absolutely improves the overall service effectiveness. Bi-directional feedback between DOTS agents is required for the increased awareness of each party, supporting superior and highly efficient attack mitigation service.

4. DOTS Telemetry Attributes

This section outlines the set of DOTS telemetry attributes. The ultimate objective of these attributes is to allow for the complete knowledge of attacks and the various particulars that can best characterize attacks.

The description and motivation behind each attribute were presented in Section 3. DOTS telemetry attributes are optionally signaled and therefore MUST NOT be treated as mandatory fields in the DOTS signal channel protocol.

4.1. Pre-mitigation DOTS Telemetry Attributes

The following pre-mitigation telemetry attributes can be signaled from the DOTS client to the DOTS server.

- o DISCUSSION NOTES: (1) Some telemetry can be communicated using DOTS data channel. (2) Evaluate the risk of fragmentation, or (3) check if we can define a dedicated URI for telemetry (e.g.: use ./telemetry). Some of the information is not specific to each mitigation request.

4.1.1. Total Traffic Normal Baseline

The low percentile (10th percentile), mid percentile (50th percentile), high percentile (90th percentile) and peak values of "Total traffic normal baselines" measured in kilobytes per second or megabytes per second or gigabytes per second.

4.1.2. Total Pipe Capability

The limit of traffic volume, in kilobytes per second or megabytes per second or gigabytes per second. This attribute represents the DOTS client domain pipe limit.

- o NOTE: Multi-homing case to be considered.

4.1.3. Total Attack Traffic

The low percentile (10th percentile), mid percentile (50th percentile), high percentile (90th percentile) and peak values of total attack traffic measured in kilobytes per second or megabytes per second or gigabytes per second.

4.1.4. Total Traffic

The low percentile (10th percentile), mid percentile (50th percentile), high percentile (90th percentile) and peak values of total traffic during a DDoS attack measured in kilobytes per second or megabytes per second or gigabytes per second.

4.1.5. Attack Details

Various information and details that describe the on-going attacks that needs to be mitigated by the DOTS server. The attack details need to cover well-known and common attacks (such as a SYN Flood) along with new emerging or vendor-specific attacks. The following fields describing the on-going attack are discussed:

vendor-id: Vendor ID is a security vendor's Enterprise Number as registered with IANA [Enterprise-Numbers]. It is a four-byte integer value.

This is a mandatory sub-attribute.

attack-id: Unique identifier assigned by the vendor for the attack.

This is a mandatory sub-attribute.

attack-name: Textual representation of attack description. Natural Language Processing (e.g., word embedding) can possibly be used to map the attack description to an attack type. Textual representation of attack solves two problems (a) avoids the need to create mapping tables manually between vendors (2) Avoids the need to standardize attack types which keep evolving.

This is a mandatory sub-attribute

attack-severity: Attack severity. Emergency (0), critical (1) and alert (2).

This is an optional sub-attribute

4.2. DOTS Client to Server Mitigation Efficacy DOTS Telemetry Attributes

The mitigation efficacy telemetry attributes can be signaled from the DOTS client to the DOTS server as part of the periodic mitigation efficacy updates to the server.

4.2.1. Total Attack Traffic

The low percentile (10th percentile), mid percentile (50th percentile), high percentile (90th percentile) and peak values of total attack traffic the DOTS client still sees during the active mitigation service measured in kilobytes per second or megabytes per second or gigabytes per second.

4.2.2. Attack Details

The overall attack details as observed from the DOTS client perspective during the active mitigation service. The same attributes defined in Section 4.1.5 are applicable here.

4.3. DOTS Server to Client Mitigation Status DOTS Telemetry Attributes

The mitigation status telemetry attributes can be signaled from the DOTS server to the DOTS client as part of the periodic mitigation status update.

4.3.1. Mitigation Status

As defined in [RFC8612], the actual mitigation activities can include several countermeasure mechanisms. The DOTS server SHOULD signal the current operational status to each relevant countermeasure. A list of attacks detected by each countermeasure. The same attributes defined for Section 4.1.5 are applicable here for describing the attacks detected and mitigated.

5. DOTS Telemetry YANG Module

5.1. Tree Structure

TODO

5.2. YANG Module

TODO

6. IANA Considerations

6.1. DOTS Signal Channel CBOR Mappings Registry

This specification registers the DOTS telemetry attributes in the IANA "DOTS Signal Channel CBOR Mappings" registry established by [I-D.ietf-dots-signal-channel].

The DOTS telemetry attributes defined in this specification are comprehension-optional parameters.

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

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
TODO				

6.2. DOTS Signal Telemetry 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:TODO
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 [RFC7950] within the "YANG Parameters" registry.

```
name: ietf-dots-telemetry
namespace: urn:ietf:params:xml:ns:yang:TODO
maintained by IANA: N
prefix: dots-telemetry
reference: RFC XXXX
```

7. Security Considerations

Security considerations in [I-D.ietf-dots-signal-channel] need to be taken into consideration.

8. Acknowledgements

The authors would like to thank Flemming Andreassen, Liang Xia, and Kaname Nishizuka co-authors of <https://tools.ietf.org/html/draft-doron-dots-telemetry-00> draft and everyone who had contributed to that document.

9. References

9.1. Normative References

- [Enterprise-Numbers]
"Private Enterprise Numbers", 2005, <<http://www.iana.org/assignments/enterprise-numbers.html>>.
- [I-D.ietf-dots-data-channel]
Boucadair, M. and R. K, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", draft-ietf-dots-data-channel-30 (work in progress), July 2019.
- [I-D.ietf-dots-signal-channel]
K, R., Boucadair, M., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", draft-ietf-dots-signal-channel-35 (work in progress), July 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative References

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

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

Email: kondtir@gmail.com

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Ehud Doron
Radware Ltd.
Raoul Wallenberg Street
Tel-Aviv 69710
Israel

Email: ehudd@radware.com