

DOTS WG
Internet-Draft
Intended status: Informational
Expires: November 9, 2017

R. Dobbins, Ed.
Arbor Networks
S. Fouant

D. Migault
Ericsson
R. Moskowitz
HTT Consulting
N. Teague
Verisign Inc
L. Xia
Huawei
K. Nishizuka
NTT Communications
May 8, 2017

**Use cases for DDoS Open Threat Signaling (DDoS) Open Threat Signaling
draft-ietf-dots-use-cases-05**

Abstract

The DDoS Open Threat Signaling (DOTS) effort is intended to provide a dynamic solution for DDoS cooperation between networks to appropriately react to DDoS attacks. This document presents use cases to evaluate the interactions expected between the DOTS components as well as the DOTS exchanges. The purpose of the use cases is to identify the interacting DOTS components, how they collaborate and what are the types of information to be exchanged.

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 <http://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 9, 2017.

Copyright Notice

Copyright (c) 2017 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 (<http://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 and Acronyms	3
3.	Use Cases Scenarios	3
3.1.	Inter-domain Use Cases	4
3.1.1.	Enterprise with an upstream transit provider DDoS mitigation Service	4
3.1.2.	Enterprise with a Cloud DDoS Mitigation Provider	5
3.2.	Intra-domain Use Cases	6
3.2.1.	Homenet DDoS Detection Collaboration for ISP network management	6
3.2.2.	DDoS Orchestration	9
4.	Security Considerations	12
5.	IANA Considerations	12
6.	Acknowledgments	12
7.	Informative References	12
	Authors' Addresses	13

[1.](#) Introduction

Currently, distributed denial-of-service (DDoS) attack mitigation solutions are largely based upon siloed, proprietary communications schemes which result in vendor lock-in. As a side-effect, this makes the configuration, provisioning, operation, and activation of these solutions a highly manual and often time-consuming process. Additionally, coordination of multiple DDoS mitigation solutions simultaneously engaged in defending the same organization (resources) against DDoS attacks is fraught with both technical and process-related hurdles. This greatly increase operational complexity and often results in suboptimal DDoS attack mitigation efficacy.

The DDoS Open Threat Signaling (DOTS) effort is intended to specify a protocol that facilitates interoperability between multivendor DDoS mitigation solutions and ensures more automation in term of mitigation requests and attack characterization patterns. As DDoS solutions are broadly heterogeneous among different vendors, the primary goal for DOTS is to provide a high level interaction with these DDoS solutions such as initiating or terminating DDoS mitigation assistance.

It should be noted that DOTS is not in and of itself intended to perform orchestration functions duplicative of the functionality being developed by the [[I2NSF](#)] WG; rather, DOTS is intended to allow devices, services, and applications to request DDoS attack mitigation assistance and receive mitigation status updates.

These use cases are expected to provide inputs for the design of the DOTS protocol(s).

2. Terminology and Acronyms

This document makes use of the terms defined in [[I-D.ietf-dots-requirements](#)].

In addition, this document introduces the following terms:

Inter-domain: a DOTS communications relationship between distinct organizations with separate spans of administrative control. Typical inter-domain DOTS communication relationships would be between a DDoS mitigation service provider and an end-customer who requires DDoS mitigation assistance; between multiple DDoS mitigation service providers coordinating mutual defense of a mutual end-customer; or between DDoS mitigation service providers which are requesting additional DDoS mitigation assistance in for attacks which exceed their inherent DDoS mitigation capacities and/or capabilities.

Intra- domain: a DOTS communications relationship between various (network) elements that are owned and operated by the same administrative entity. A typical intra-domain DOTS communications relationship would be between DOTS agents [[I-D.ietf-dots-requirements](#)] within the same organization.

3. Use Cases Scenarios

This section provides a high-level description of scenarios addressed by DOTS. In both sub-sections, the scenarios are provided in order to illustrate the use of DOTS in typical DDoS attack scenarios. They

are not definitive, and other use cases are expected to emerge with widespread DOTS deployment.

All scenarios present a coordination between the targeted organization, the DDoS attack telemetry and the mitigator. The coordination and communication between these entities depends, for example, on the characteristic or functionality of the entity itself, the reliability of the information provided by DDoS attack telemetry, and the business relationship between the DDoS target domain and the mitigator.

More explicitly, in some cases, the DDoS attack telemetry may simply activate a DDoS mitigation, whereas in other cases, it may collaborate by providing some information about an attack. In some cases, the DDoS mitigation may be orchestrated, which includes selecting a specific appliance as well as starting/ending a mitigation.

3.1. Inter-domain Use Cases

3.1.1. Enterprise with an upstream transit provider DDoS mitigation Service

In this scenario, an enterprise network with self-hosted Internet-facing properties such as Web servers, authoritative DNS servers, and VoIP service platforms has a DDoS mitigation system (DMS) deployed to protect those servers, applications, and network resources from DDoS attacks. In addition to their on-premise DDoS defense capability, they have contracted with their Internet access provider for DDoS mitigation services which threaten to overwhelm their WAN interconnection link(s) bandwidth.

The DMS is configured such that if the incoming Internet traffic volume exceeds, e.g., 50% of the provisioned upstream Internet interconnection link(s) capacity, the DMS will request DDoS mitigation assistance from the upstream access provider.

Before any communication takes place between DOTS agents, security credentials are provisioned on these agents so that only authorized entities can trigger mitigation actions.

The communication to trigger, manage, and terminate a DDoS mitigation between the enterprise DMS and the access provider(s) is performed using DOTS. The enterprise DMS implements a DOTS client while the access provider implements a DOTS server. A DOTS client can establish communications with multiple DOTS servers, if the enterprise is multi-homed or of distinct access technologies are used (e.g., fixed, LTE).

When the DMS detects an inbound DDoS attack targeting the enterprise resources, it immediately begins mitigating the attack.

During the course of the attack, the inbound traffic volume exceeds the 50% threshold; the DMS DOTS client signals its DOTS server(s) on the upstream access provider network(s) to initiate DDoS mitigation immediately. The DOTS server signals the DOTS client that it can service this request, and mitigation is initiated on the access provider network.

Over the course of the attack, the DOTS server on the transit provider network periodically signals the DOTS client on the enterprise DMS in order to provide mitigation status information, statistics related to DDoS attack traffic mitigation, and related information. Such information are collected by the DOTS server, but the way these are collected are outside of DOTS. Once the DDoS attack has ended, the DOTS server signals the enterprise DMS DOTS client that the attack has subsided. This signal may not be sent immediately, but once the peace time is judged stable; the duration observation of the peace time after an attack is deployment-specific.

The enterprise DMS then requests that DDoS mitigation services on the upstream access provider network be terminated. The DOTS server on the access provider network receives this request, communicates with the access provider orchestration system controlling its DDoS mitigation system to terminate attack mitigation, and once the mitigation has ended, confirms the end of upstream DDoS mitigation service to the enterprise DMS DOTS client.

Request termination will be repeated with each of the upstream DOTS servers reachable through links that were under DDoS attack.

3.1.2. Enterprise with a Cloud DDoS Mitigation Provider

This use case details an enterprise that has a local DDoS detection and classification capability and may or may not have a mitigation capability. The enterprise is contracted with a cloud DDoS mitigation provider who can redirect (off-ramp) traffic away from the enterprise, provide scrubbing services, and return "clean" traffic back to the enterprise (on-ramp) on an ad-hoc, on demand basis.

The enterprise may, either by hard coding or on a case by case basis, determine thresholds at which a request for mitigation is triggered indicating to the cloud provider that traffic should be redirected and scrubbed.

The communication to trigger, manage, and terminate a DDoS mitigation between the enterprise and the Cloud provider is performed using

DOTS. The enterprise implements a DOTS client while the Cloud Provider implements a DOTS server.

The enterprise detection and classification systems encompass a DOTS client and the cloud provider a DOTS server.

When an attack is detected an automated or manual DOTS mitigation request will be generated and sent to the cloud provider. The cloud provider will assess the request for validity and if passed a mitigation action may then be initiated. This action will usually involve the off-ramp of all traffic destined to the target for further scrutiny and filtering by the cloud provider. This should not only result in an alleviation of pressure on the enterprise network but also on its upstream provider and peers. How traffic redirection is implemented is out of scope.

The cloud provider should signal via DOTS to the enterprise that a mitigation request has been received and acted upon and should also include a basic situational status of the attack. The cloud provider may respond periodically with additional updates on the status to enable the enterprise to make an informed decision on whether to maintain or cancel the mitigation. An alternative approach would be for the DOTS client mitigation request to include a time to live (TTL) for the mitigation which may be extended by the client should the attack still be ongoing as the TTL reaches expiration.

A variation of this use case may be that the enterprise is providing a flow-based monitoring and analysis service to customers whose networks may be protected by any one of a number of 3rd party providers. The enterprise in question may integrate with these 3rd party providers using DOTS and signal accordingly when a customer is attacked - the enterprise may then manage the life-cycle of the attack on behalf of the enterprise.

3.2. Intra-domain Use Cases

3.2.1. Homenet DDoS Detection Collaboration for ISP network management

Home networks run with (limited) bandwidth as well as limited routing resources, while they are expected to provide services reachable from the outside [[RFC7368](#)]. This makes such networks some easy targets to DDoS attacks via their WAN interface. As these DDoS attacks are easy to perform, they may remain undetected by the upstream ISP. When the CPE is congested, the customer is likely to call the ISP hotline. In order to improve the quality of experience of the connectivity as well as to automate the request for DDoS mitigation, ISPs are likely to consider a standard mean for CPEs to automatically inform a

dedicated service mitigation platform when they are under a suspected DDoS.

Note also that this section only considers DDoS attacks CPE or services in the home network are encountering. This differs from DDoS attacks the CPE or any device within the home network may take part of - such as botnets. In the later attacks, the home network generates traffic under the control of a botmaster. Such attacks may only be detected once the attacks have been characterized. It would be tempting to consider a feature in the DOTS protocol to allow a DOTS server to inform a CPE that some suspect traffic is being sent by the CPE so that appropriate actions are undertaken by the CPE/user. Nevertheless, this feature would require some interaction with the CPE administrator. Such scenario is outside the scope of this document.

In this use case, ISPs are willing to prevent their customer undergoing DDoS attacks in order to enhance the quality of experience of their customers, to avoid unnecessary costly call on hot lines as well as to optimize the bandwidth of their network. A key challenge for the ISP is to detect DDoS attacks. In fact, DDoS detection is not only fine grained but is also expected to be different for each home network or small businesses networks (SOHO), and the ISP is unlikely to have sufficient resource to inspect the traffic of all its customers.

In order to address these challenges, ISPs are delegating the DDoS detection to CPE of home network or SOHO. Outsourcing the detection on the CPE provides the following advantages to the ISP: 1) Avoid the ISP to dedicate a huge amount of resource for deep packet inspection over a large amount of traffic with a specific security policies associated to each home network. It is expected that such traffic only constitutes a small fraction of the total traffic the ISP is responsible for. 2) DDoS detection is deployed in a scalable way. 3) Provide more deterministic DDoS attack detection. For example, what could be suspected to be an UDP flood by the ISP may be consented by the terminating point hosted in the home network or SOHO. In fact, without specific home network security policies, the ISP is likely to detect DDoS attack over regular traffic or to miss DDoS attacks targeting a specific home network or CPE. In the first case, this would result in the ISP spending unnecessary resources and in the second case this would directly impact the quality of experience of the customer.

Note that in this scenario slightly differs from the "Enterprise with an upstream transit provider DDoS mitigation Service" scenario described in [Section 3.1.1](#). In this scenario, the detection DDoS is motivated by the ISP in order to operate appropriately its network.

For that purpose, it requires some collaboration with the home network. In [Section 3.1.1](#), the target network requests a mitigation service from the upstream transit provider in order to operate its services.

Even though the motivations differ, there are still significant advantages for the home network to collaborate. On the home network or SOHO perspective such collaboration provides the following advantages: 1) If it removes the flows contributing to a DDoS attacks, then it enhances the quality of experience of the users of the targeted services or the entire home network. 2) If mitigation is being handled by the ISP rather than the home network, then it reduces the management of DDoS attacks by the network administrator which involves detection as well as mitigation as well as the provisioning of extra resources. 3) If the DDoS detection is based on information specific to the home network, such as for example the description of the services, the hosts capacities or the network topology, then performing the DDoS detection by the home network instead of the ISP avoids the home network to leak private information to the ISP. In that sense, it better preserves the home network or SOHO privacy while enabling a better detection. However, the request for mitigation may still leak some informations. ISPs must not retrieve sensitive data without the consent of the user. This is usually captured in administrative contracts that are out of scope of this document.

When the CPE suspects an attack, it notifies automatically or the ISP. The contact address of the DDoS Mitigation service of the ISP may be hard coded or may be configured manually or automatically (e.g., eventually the DHCP server may provide the DDoS mitigation service via specific DHCP options).

The communication to trigger a DDoS mitigation between the home network and the ISP is performed using DOTS. The home network CPE implements a DOTS client while the ISP implements a DOTS server.

The DOTS client on the CPE monitors the status of CPE's resource and WAN link bandwidth usage. If something unusual happens based on preconfigured throughput, traffic patten, explicit action from the user, or some heuristics methods, the DOTS client sends a DOTS mitigation request to the ISP DOTS server. Typically, a default configuration with no additional information associated to the DOTS mitigation request is expected. The ISP derives traffic to mitigate from the CPE IP address.

In some cases, the DOTS mitigation request contains options such as some IP addresses or prefixes that belongs to a whitelist or a blacklist. In this case, the white and black lists are not

associated to some analysis performed by the CPE -- as the CPE is clearly not expected to analyze such attacks. Instead these are part of some configuration parameters. For example, in the case of small business, one may indicate specific legitimate IP addresses such as those used for VPNs, or third party services the company is likely to set a session. Similarly, the CPE may provide the IP addresses targeting the assets to be protected inside the network. Note that the IP address is the IP address used to reach the asset from the internet, and as such is expected to be globally routable. Such options may include the IP address as well as a service description. Similarly to the previous blacklist and whitelist, such information are likely not derived from a traffic analysis performed by the CPE, but instead are more related to configuration parameters.

Upon receiving the DOTS mitigation request, the DOTS server acknowledges its reception and confirms DDoS mitigation starts or not. Such feed back is mostly to avoid retransmission of the request.

Note that the ISP is connected to multiple CPEs and as such the CPE can potentially perform DDoS attack to the DOTS server. ISP may use gateways to absorb the traffic. These gateways, will typically aggregate a smaller number of CPEs and retransmit to the destination DOTS Server a selected information. Note that such gateways may somehow act as a DOTS relay, which is implemented with a DOTS Server and a DOTS Client. Note also that the case of a large DDoS attack targeting simultaneously multiple CPEs is expected to be detected and mitigated by the upstream architecture, eventually without DOTS alerts sent by each single CPE.

ISP may activate mitigation for the traffic associated to the CPE sending the alert or instead to the traffic associated to all CPE. Such decisions are not part of DOTS, but instead depend on the policies of the ISP.

It is unlikely the CPE will follow the status of the mitigation. The ISP is only expected to inform the CPE the mitigation has been stopped.

Upon receipt of such notification the CPE may, for example, re-activate the monitoring jobs and thus is likely to provide some further DOTS alert.

3.2.2. DDoS Orchestration

In this use case, one or multiple DDoS telemetry systems like a flow collector monitor a network -- typically an ISP network. Upon detection of a DDoS attack, these telemetry systems alert an

orchestrator in charge of coordinating the various DDoS mitigation systems within the domain. The telemetry systems may be configured to provide some necessary or useful pieces of information, such as a preliminary analysis of the observation to the orchestrator.

The orchestrator analyses the various information it receives from specialized equipment, and elaborates one or multiple DDoS mitigation strategies. In some case, a manual confirmation may also be required to choose a proposed strategy or to start the DDoS mitigation. The DDoS mitigation may consists in multiple steps such as configuring the network, the various hardware or already instantiated DDoS mitigation functions. In some cases, some specific virtual DDoS mitigation functions need to be instantiated and properly chained between each other. Eventually, the coordination of the mitigation may involve external DDoS resources such as a transit provider ([Section 3.1](#)) or a cloud provider ([Section 3.1.2](#)).

The communication to trigger a DDoS mitigation between the telemetry and monitoring systems and the orchestrator is performed using DOTS. The telemetry systems implements a DOTS client while the Orchestrator implements a DOTS server.

The communication between a network administrator and the orchestrator is also performed using DOTS. The network administrator via its web interfaces implements a DOTS client while the Orchestrator implements a DOTS server.

The communication between the Orchestrator and the DDoS mitigation systems is performed using DOTS. The Orchestrator implements a DOTS client while the DDoS mitigation systems implement a DOTS server.

The configuration aspects of each DDoS mitigation systems, as well as the instantiations of DDoS mitigation functions or network configuration is not part of DOTS. Similarly the discovery of the available DDoS mitigation functions is not part of DOTS.

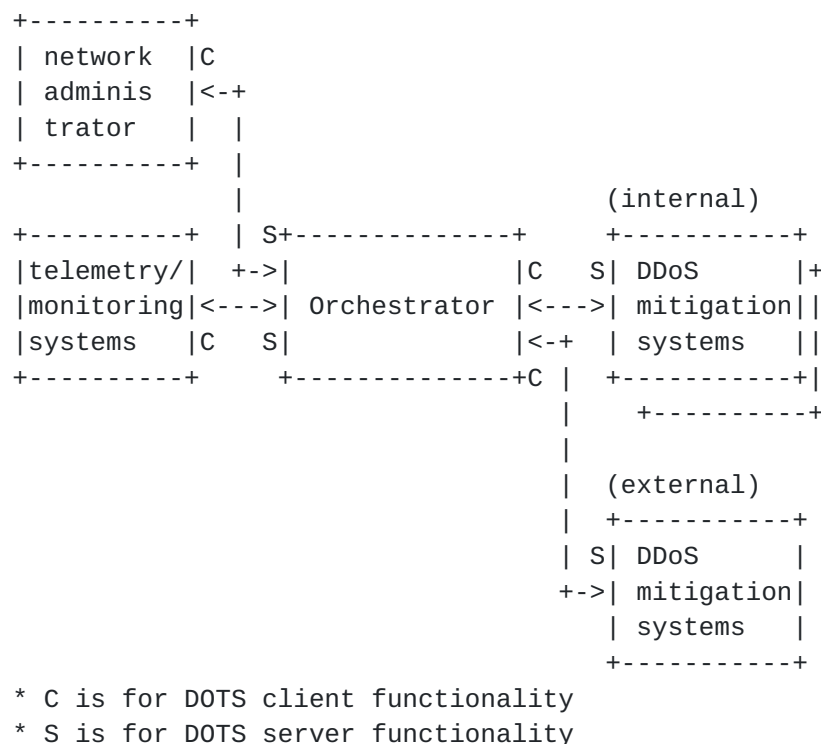


Figure 1: DDoS Orchestration

The telemetry systems monitor various traffic network and perform their measurement tasks. They are configured so that when an event or some measurements reach a predefined level to report a DOTS mitigation request to the Orchestrator. The DOTS mitigation request may be associated with some element such as specific reporting.

Upon receipt of the DOTS mitigation request from the telemetry system, the Orchestrator responds with an acknowledgement, to avoid retransmission of the request for mitigation. The status of the DDoS mitigation indicates the Orchestrator is in an analysing phase. The Orchestrator begins collecting various information from various telemetry systems in order to correlate the measurements and provide an analysis of the event. Eventually, the Orchestrator may ask additional information to the telemetry system, however, the collection of these information is performed outside DOTS.

The Orchestrator may be configured to start a DDoS mitigation upon approval from a network administrator. The analysis from the orchestrator is reported to the network administrator via a web interface. If the network administrator decides to start the mitigation, she orders through her web interface a DOTS client to send a request for DDoS mitigation. This request is expected to be associated with a context that identifies the DDoS mitigation selected.

Upon receiving the DOTS request for DDoS mitigation from the network administrator, the orchestrator orchestrates the DDoS mitigation according to the specified strategy. Its status indicates the DDoS mitigation is starting while not effective.

Orchestration of the DDoS mitigation systems works similarly as described in [Section 3.1](#) and [Section 3.1.2](#). The Orchestrator indicates with its status whether the DDoS mitigation is effective.

When the DDoS mitigation is finished on the DDoS mitigation systems, the Orchestrator indicates to the telemetry systems as well as to the network administrator the DDoS mitigation is finished.

4. Security Considerations

DOTS is at risk from three primary attacks: DOTS agent impersonation, traffic injection, and signaling blocking. Associated security requirements and additional ones are defined in [\[I-D.ietf-dots-requirements\]](#).

Impersonation and traffic injection mitigation can be managed through current secure communications best practices. DOTS is not subject to anything new in this area. One consideration could be to minimize the security technologies in use at any one time. The more needed, the greater the risk of failures coming from assumptions on one technology providing protection that it does not in the presence of another technology.

5. IANA Considerations

No IANA considerations exist for this document at this time.

6. Acknowledgments

The authors would like to thank among others Tirumaleswar Reddy, , Andrew Mortensen, Mohamed Boucadair, the DOTS WG chairs Roman D. Danyliw and Tobias Gondrom for their valuable feed backs.

7. Informative References

- [I-D.ietf-dots-requirements]
Mortensen, A., Moskowitz, R., and T. Reddy, "Distributed Denial of Service (DDoS) Open Threat Signaling Requirements", [draft-ietf-dots-requirements-04](#) (work in progress), March 2017.
- [I2NSF] "Interface to Network Security Functions (i2nsf)",
<<https://datatracker.ietf.org/wg/i2nsf/about/>>.

[RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", [RFC 7368](https://www.rfc-editor.org/info/rfc7368), DOI 10.17487/RFC7368, October 2014, <<http://www.rfc-editor.org/info/rfc7368>>.

Authors' Addresses

Roland Dobbins (editor)
Arbor Networks
30 Raffles Place
Level 17 Chevron House
Singapore 048622
Singapore

Email: rdobbins@arbor.net

Stefan Fouant

Email: stefan.fouant@copperriverit.com

Daniel Migault
Ericsson
8400 boulevard Decarie
Montreal, QC H4P 2N2
Canada

Phone: +1 514-452-2160

Email: daniel.migault@ericsson.com

Robert Moskowitz
HTT Consulting
Oak Park, MI 48237
USA

Email: rgm@labs.htt-consult.com

Nik Teague
Verisign Inc
12061 Bluemont Way
Reston, VA 20190
USA

Phone: +44 791 763 5384

Email: nteague@verisign.com

Liang Xia
Huawei
No. 101, Software Avenue, Yuhuatai District
Nanjing
China

Email: Frank.xialiang@huawei.com

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

Email: kaname@nttv6.jp

