

DOTS WG
Internet-Draft
Intended status: Informational
Expires: September 28, 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
March 27, 2017

Use cases for DDoS Open Threat Signaling
draft-ietf-dots-use-cases-04.txt

Abstract

The DDoS Open Threat Signaling (DOTS) effort is intended to provide a protocol that facilitates interoperability between multivendor solutions/services. 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 component, 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 September 28, 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
2.1.	Requirements Terminology	3
2.2.	Acronyms	3
2.3.	Terms	3
3.	Use Cases Scenarios	4
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 on Cloud DDoS mitigation provider	5
3.2.	Intra-domain Use Cases	6
3.2.1.	Homenet DDoS protection by ISP	6
3.2.2.	DDoS Orchestration	8
4.	Security Considerations	10
5.	IANA Considerations	10
6.	Acknowledgments	10
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	11
	Authors' Addresses	11

[1.](#) Introduction

Currently, distributed denial-of-service (DDoS) attack mitigation solutions/services are largely based upon siloed, proprietary communications paradigms which result in vendor/service 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/services simultaneously engaged in defending the same organization 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 provide a protocol that facilitates interoperability between multivendor DDoS mitigation solutions/services. As DDoS solutions/services are broadly heterogeneous among different vendors, the primary goal for DOTS is to provide a high level interaction with these DDoS solutions/services 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 from systems of this nature.

The use cases presented in the document are intended to provide examples of communications interactions DOTS-enabled nodes in both inter- and intra-organizational DDoS mitigation scenarios. These use cases are expected to provide inputs for the design of the DOTS protocol(s).

2. Terminology and Acronyms

2.1. Requirements Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2.2. Acronyms

This document makes use of the same terminology and definitions as [[I-D.ietf-dots-requirements](#)], except where noted.

2.3. Terms

Inter-organizational: a DOTS communications relationship between distinct organizations with separate spans of administrative control. Typical inter-organizational DOTS communication relationships would be between a DDoS mitigation service provider and an end-customer organizational which 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-organizational: a DOTS communications relationship between various elements within a single span of administrative control. A typical intra-organizational DOTS communications relationship would be between DOTS clients, DOTS gateways, and DOTS servers within the same organization.

3. Use Cases Scenarios

This section provides a high-level description of scenarios addressed by DOTS. In both 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 entity depends, for example on the characteristic or functionality of the equipment, 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 PBXes has an intelligent DDoS mitigation system (IDMS) deployed to protect those servers and applications from DDoS attacks. In addition to their on-premise DDoS defense capability, they have contracted with their Internet transit provider for DDoS mitigation services which threaten to overwhelm their transit link bandwidth.

The IDMS is configured such that if the incoming Internet traffic volume exceeds 50% of the provisioned upstream Internet transit link

capacity, the IDMS will request DDoS mitigation assistance from the upstream transit provider.

The communication to trigger, manage, and finalize a DDoS mitigation between the enterprise IDMS and the transit provider is performed using DOTS. The enterprise IDMS implements a DOTS client while the transit provider implements a DOTS server.

When the IDMS detects an inbound DDoS attack targeting the enterprise servers and applications, it immediately begins mitigating the attack.

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

Over the course of the attack, the DOTS server on the transit provider network periodically signals the DOTS client on the enterprise IDMS in order to provide mitigation status information, statistics related to DDoS attack traffic mitigation, and related information. Once the DDoS attack has ended, the DOTS server signals the enterprise IDMS DOTS client that the attack has subsided.

The enterprise IDMS then requests that DDoS mitigation services on the upstream transit provider network be terminated. The DOTS server on the transit provider network receives this request, communicates with the transit 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 IDMS DOTS client.

3.1.2. Enterprise with on 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 (offramp) traffic away from the enterprise, provide scrubbing services and return clean traffic back to the enterprise (onramp) 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 finalize 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.

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 protection by ISP

In this use case home networks or small businesses networks (SOHO), subscribe with their upstream ISP a DDoS mitigation service.

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 organizations some easy targets to DDoS attacks. In addition, these DDoS attacks might even not be noticed by the upstream ISP.

This scenario is considered as an intra-domain as ISPs have a specific relationship with these customers. The ISP is the connectivity provider, and in some cases, they even provides the CPE with a set of associated services. Moreover, in case of any connectivity issue the customer is likely to call the hotline. In order to improve the QoS of the connectivity as well as to automate the request for DDoS mitigation, ISP is likely to consider a standard mean for CPEs to notify when they are under a suspected DDoS. Such notification may be triggered automatically or manually. As the ISP and the customer share a common interest in mitigating the DDoS attack, this slightly differs from cases where a contract is negotiated with a third party, such as in the inter-domain use cases.

In most cases, CPEs are unlikely to diagnose whether an DDoS attack is ongoing or not and simply rely on the upstream equipment provided by the ISP for detection and potential mitigation.

The DDoS Mitigation service of the ISP may be hard coded or may be configured by the customer manually or automatically while the CPE is being connected to the Internet -- 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 link bandwidth usage. If something unusual happens based on preconfigured throughput 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 respectively to 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 provides the IP addresses of the assets to be protected inside the network. Such options may include the IP address as well as a service description. Similarly to the previous blacklist and whitelist, such information are 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 ISP 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 relays to absorb the traffic. In addition, such attack may be triggered by a large scale DDoS attack, which is expected to be detected and mitigated by the upstream architecture.

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 network administrator.

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 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 telemetry systems or monitoring devices 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 informations, such as a preliminary analysis of the observation to the orchestrator.

The orchestrator analyses the various information it receives from specialized equipments, and elaborates one or multiple DDoS mitigation strategies. In some case, a manual confirmation may also be required to chose 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 involved 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 to select a DDoS strategy by 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 pat of DOTS.

The Telemetry or monitoring systems monitors each various traffic network and each performs their measurement tasks. They are configure 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, or analysis.

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 informations from various telemetry systems on the network in order to correlate the measurements and provide an analyse of the event. Eventually, the orchestrator may ask additional informations to the telemetry system that just sent the DOTS request, 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 order 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. It status first indicates the DDoS mitigation is starting while not effective. In fact the

orchestrator is expected to proceed to a significant number of configurations.

Orchestration of the DDoS mitigation systems works similarly as described in [Section 3.1](#) or [Section 3.1.2](#). The orchestrator indicates with its status 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. The DOTS protocol MUST be designed for minimal data transfer to address the blocking risk.

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.

Additional details of DOTS security requirements may be found in [\[I-D.ietf-dots-requirements\]](#).

5. IANA Considerations

No IANA considerations exist for this document at this time.

6. Acknowledgments

TBD

7. References

7.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.

7.2. Informative References

- [APACHE] "Apache mod_security", <<https://www.modsecurity.org>>.
- [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.
- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", [BCP 165](#), [RFC 6335](#), DOI 10.17487/RFC6335, August 2011, <<http://www.rfc-editor.org/info/rfc6335>>.
- [RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", [RFC 7368](#), DOI 10.17487/RFC7368, October 2014, <<http://www.rfc-editor.org/info/rfc7368>>.
- [RRL] "BIND RRL", <<https://deephought.isc.org/article/AA-00994/0/Using-the-Response-Rate-Limiting-Feature-in-BIND-9.10.html>>.

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

