DOTS
Internet-Draft
Intended status: Informational
Expires: September 29, 2018

R. Dobbins
Arbor Networks
D. Migault
Ericsson
S. Fouant

R. Moskowitz
HTT Consulting
N. Teague
Verisign
L. Xia
Huawei
K. Nishizuka
NTT Communications
March 28, 2018

Use cases for DDoS Open Threat Signaling
draft-ietf-dots-use-cases-11

Abstract

The DDoS Open Threat Signaling (DOTS) effort is intended to provide a protocol to facilitate interoperability across disparate DDoS mitigation solutions and services. This document presents use cases which describe the interactions expected between the DOTS components as well as DOTS messaging exchanges. The purpose of describing use cases is to identify the interacting DOTS components, how they collaborate and what are the types of information to be exchanged.

Table of Contents

1. Introduction

At the time of writing, distributed denial-of-service (DDoS) attack mitigation solutions are largely based upon siloed, proprietary communications schemes with vendor lock-in as a side-effect. This can result in the configuration, provisioning, operation, and activation of these solutions being a highly manual and often time-consuming process. Additionally, coordination of multiple DDoS mitigation solutions simultaneously is fraught with both technical and process-related hurdles. This greatly increases operational complexity which, in turn, can degrade the efficacy of mitigations.

The DDoS Open Threat Signaling (DOTS) effort is intended to specify a protocol that facilitates interoperability between diverse DDoS mitigation solutions and ensures greater integration in term of mitigation requests and attack characterization patterns. As DDoS solutions are broadly heterogeneous among vendors, the primary goal of DOTS is to provide high-level interaction amongst differing DDoS solutions, such as initiating, terminating DDoS mitigation assistance or requesting the status of a DDoS mitigation.

This document provides use cases to provide inputs for the design of the DOTS protocol(s) as well as to illustrate the purpose of goals. The use cases are not exhaustive and future use cases are expected to emerge as DOTS is adopted and evolves.

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 [RFC2119].

2.2. Acronyms

This document makes use of the same terminology and definitions as [I-D.ietf-dots-requirements]. In addition it uses the terms defined below:

- DDoS Mitigation Service Provider: designates the administrative entity providing the DDoS Mitigation Service.

- DDoS Mitigation Service: designates a service provides to a customer. Services usually involves Service Level Agreement (SLA) that have to be met. It is the responsibility of the service provider to instantiate the DDoS Mitigation System to meet these SLA.

- DDoS Mitigation System (DMS): A system that performs DDoS mitigation. The DDoS Mitigation System may be composed by a cluster of hardware and/or software resources, but could also involve an orchestrator that may take decisions such as outsourcing partial or more of the mitigation to another DDoS Mitigation System.

- DDoS Mitigation: The action performed by the DDoS Mitigation System.

- Internet Transit Provider (ITP):

3. Use Cases

3.1. Upstream DDoS Mitigation between an Enterprise Network and an Upstream Internet Transit Provider

This use case describes how an enterprise network may take advantage of a pre-existing relation with its Internet Transit Provider (ITP) in order to mitigate a DDoS attack targeting its network. As the ITP provides connectivity to the enterprise network, it is already on the path of the inbound or outbound traffic of the enterprise network and well aware of the networking parameters associated to the enterprise network connectivity. This eases both the configuration and the instantiation of a DDoS Mitigation Service. This section considers two kind of DDoS Mitigation Service between an enterprise network and an ITP:

- The upstream ITP may instantiate a DDoS Mitigation System (DMS) upon receiving a request from the enterprise network. This typically corresponds to the case when the enterprise network is under attack.

- On the other hand, the ITP may identify an enterprise network as the source of an attack and send a mitigation request for the enterprise to mitigate this at the source.

In the first scenario, as depicted in Figure 1, an enterprise network with self-hosted Internet-facing properties such as Web servers, authoritative DNS servers, and VoIP PBXes has a DMS 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.

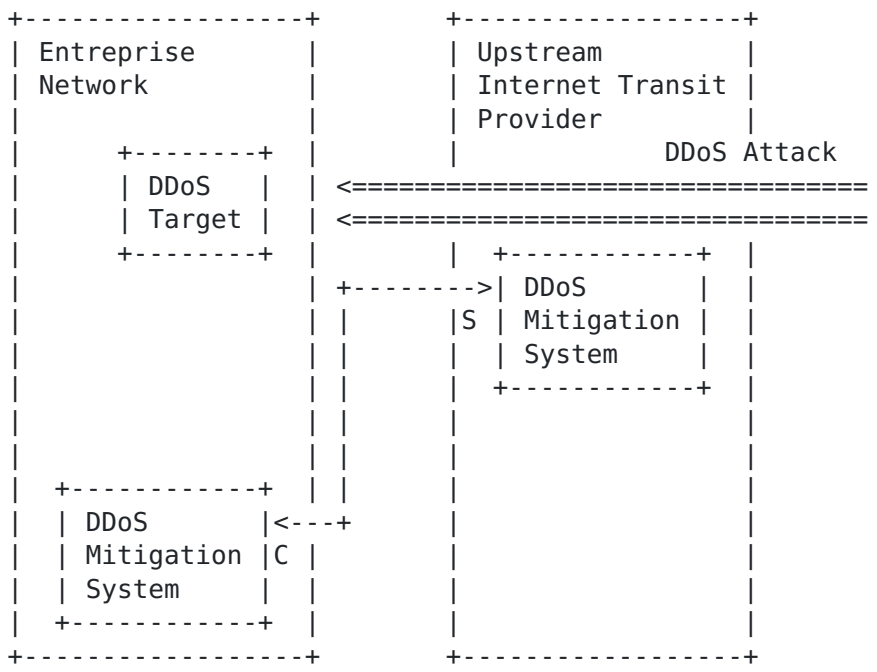

* C is for DOTS Client functionality
* S is for DOTS Server functionality

Figure 1: Upstream Internet Transit Provider DDoS Mitigation

The enterprise DMS is configured such that if the incoming Internet traffic volume exceeds 50% of the provisioned upstream Internet transit link capacity, the DMS will request DDoS mitigation assistance from the upstream transit provider.

The requests to trigger, manage, and finalize a DDoS Mitigation between the enterprise DMS and the ITP is performed using DOTS. The enterprise DMS implements a DOTS Client while the ITP implements a DOTS Server which is integrated with their DMS.

When the enterprise DMS detects an inbound DDoS attack targeting its servers and applications, it immediately begins a DDoS Mitigation.

During the course of the attack, the inbound traffic volume exceeds the 50% threshold; the DMS DOTS Client signals the DOTS Server on the upstream ITP to initiate DDoS Mitigation. The DOTS Server signals the DOTS Client that it can serve this request, and mitigation is initiated on the ITP network by the DMS.

Over the course of the attack, the DOTS Server of the ITP periodically informs the DOTS Client on the enterprise DMS mitigation status, statistics related to DDoS attack traffic mitigation, and

related information. Once the DDoS attack has ended, the DOTS Server signals the enterprise DMS DOTS Client that the attack has subsided.

The enterprise DMS then requests the ITP to terminate the DDoS Mitigation. The DOTS Server on the ITP receives this request and once the mitigation has ended, confirms the end of upstream DDoS Mitigation to the enterprise DMS DOTS Client.

The following is an overview of the DOTS communication model for this use-case:

- (a) A DDoS attack is initiated against online properties of an network organization which has deployed a DOTS-Client-capable DMS.
- (b) The DMS detects, classifies, and begins the DDoS Mitigation.
- (c) The DMS determines that its capacity and/or capability to mitigate the DDoS attack is insufficient, and sends via its DOTS Client a DOTS DDoS Mitigation request to one or more DOTS Servers residing on the upstream ITP.
- (d) The DOTS Server which receive the DOTS Mitigation request determines that they have been configured to honor requests from the requesting DOTS Client, and honored its DDoS Mitigation by orchestrating its DMS.
- (e) While the DDoS Mitigation is active, the DOTS Servers regularly transmit DOTS DDoS Mitigation status updates to the DOTS Client.
- (f) The DOTS Client transmits a DOTS DDoS Mitigation termination request to the DOTS Server.
- (g) The DOTS Server terminates DDoS Mitigation.

Note that communications between the enterprise DOTS Client and the upstream transit provider DOTS Server may take place in-band within the main Internet transit link between the enterprise and the ITP; out-of-band via a separate, dedicated wireline network link utilized solely for DOTS signaling; or out-of-band via some other form of network connectivity such as a third-party wireless 4G network link.

Note also that the DOTS Clients that sends the DOTS Mitigation request may be also triggered by a network admin that manually confirms the request to the upstream ITP, in which case the request my be sent from an application such as a web browser in a mobile phone.

Note also that when the enterprise is multihomed and connected to multiple upstream ITP, each ITP is only able to provide a DDoS Mitigation Service for the traffic it transits. As a result, the enterprise network may require to coordinate the various DDoS Mitigation Services associated to each link.

The current scenario describes the case where the DDoS Target is in the enterprise network while the DMS is provided by the upstream ITP. An alternate use case may consider the case where the ITP informs the enterprise network it is involved into an ongoing attack or that infected machines have been identified. In this case the DOTS Client and DOTS Server roles are inverted. The DOTS Client is located in the ITP network and the DOTS Server is hosted in the enterprise network. The enterprise network is then responsible to perform the DDoS Mitigation. In some case the DDoS Mitigation may be delegated back to the upstream ITP, as described in this section.

3.2. DDoS Mitigation between an Enterprise Network and third party DDoS Mitigation Service Provider

This use case differs from the previous use case in that the DDoS Mitigation Service is not provided by an upstream ITP. In other words, as represented in figure 2, the traffic does not not go through the DDoS Mitigation Service Provider by default. In order to steer the traffic to the DDoS Mitigation Service Provider, some network configuration are required. As such it may be reserved for large enterprises or large data centers.

We follow the terminology of section Section 3.1, however the Enterprise Network is not limited to the network hosting the DDoS Target. In fact, it could also be a DDoS Mitigation Service Provider that has reached it resources capacities and delegate the DDoS Mitigation to other DDoS Mitigation Service Providers, thus forming an overlay of DMS.

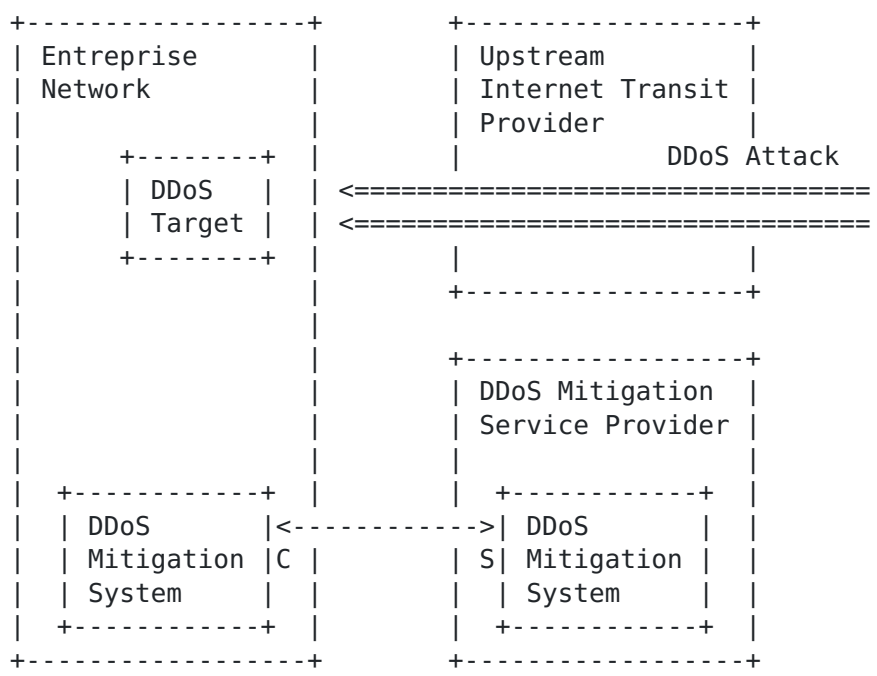

* C is for DOTS Client functionality
* S is for DOTS Server functionality

Figure 2: DDoS Mitigation between an Enterprise Network and third party DDoS Mitigation Service Provider

In this scenario, an Enterprise Network has entered into a pre-arranged DDoS mitigation assistance agreement with one or more other DDoS Mitigation Service Providers in order to ensure that sufficient DDoS mitigation capacity and/or capabilities may be activated in the event that a given DDoS attack threatens to overwhelm the ability of a given DMS to mitigate the attack on its own.

The pre-arrangement typically includes the agreement on the mechanisms used to redirect the traffic to the DDoS Mitigation Service Provider, as well as the mechanism to to re-inject the traffic back to the Enterprise Network. Redirection to the DDoS Mitigation Service Provider typically involves BGP prefix announcement eventually combined with DNS redirection, while re-injection may be performed via tunneling mechanisms such as GRE for example. Of course, such mechanisms needs to be regularly tested and evaluated.

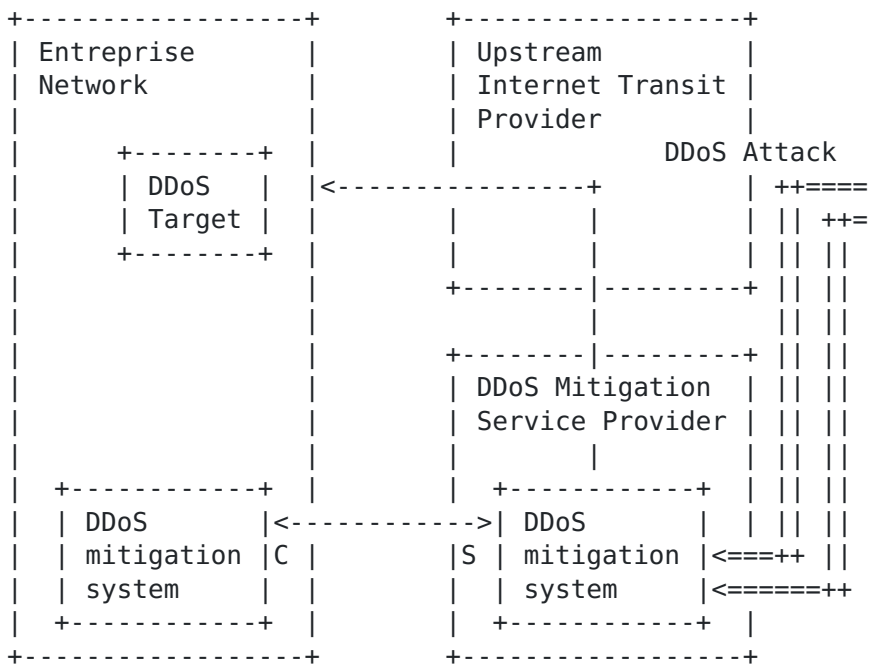

Figure 3: Redirection to a DDoS Mitigation Service Provider

When the Enterprise Network is under attack or at least is reaching its capacity or ability to mitigate a given DDoS attack traffic, the DOTS Client sends a DOTS request to the DDoS Mitigation Service Provider to initiate network traffic diversion - as represented in figure 3 - and DDoS mitigation activities. Ongoing attack and mitigation status messages may be passed between the Enterprise Network and the DDoS Mitigation Service Provider.

Once the requesting Enterprise Network is confident that the DDoS attack has either ceased or has fallen to levels of traffic/ complexity which they can handle on their own or that it has received a DOTS DDoS Mitigation termination request from a downstream Enterprise Network or DDoS Mitigation Service Provider, the requesting Enterprise Network DOTS Client sends a DOTS DDoS Mitigation termination requests to the DDoS Mitigation Service Provider.

3.3. DDoS Orchestration

In this use case, one or more DDoS telemetry systems or monitoring devices such as a flow telemetry 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 DMS within the domain. The telemetry systems may be configured to provide necessary and 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 equipements, 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 initiate a DDoS Mitigation. The DDoS Mitigation may consist of multiple steps such as configuring the network, various hardware, or updating already instantiated DDoS mitigation functions. In some cases, some specific virtual DDoS mitigation functions must be instantiated and properly ordered. Eventually, the coordination of the mitigation may involve external DDoS resources such as a transit provider or a DDoS Mitigation Service Provider.

The communications used 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 system, as well as the instantiations of DDoS mitigation functions or network configuration is not part of DOTS. Similarly, the discovery of available DDoS mitigation functions is not part of DOTS.

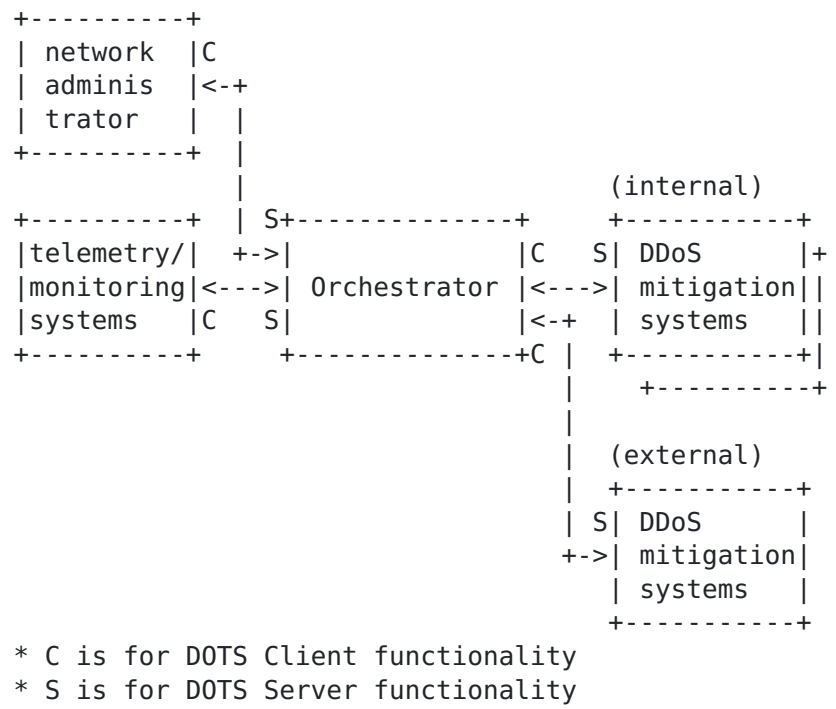

Figure 4: DDoS Orchestration

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 acknowledgment, to avoid retransmission of the request for mitigation. The status of the DDoS mitigation indicates the Orchestrator is in an analyzing 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 XXX. 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. 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

The authors would like to thank among others Tirumaleswar Reddy; Andrew Mortensen; Mohamed Boucadaire; Artyom Gavrichenkov; and the DOTS WG chairs, Roman D. Danyliw and Tobias Gondrom, for their valuable feedback.

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

7.2. 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-14 (work in progress), February 2018.

Authors' Addresses

Roland Dobbins
Arbor Networks
Singapore

EMail: rdobbins@arbor.net

Daniel Migault
Ericsson
8275 Trans Canada Route
Saint Laurent, QC 4S 0B6
Canada

EMail: daniel.migault@ericsson.com

Stefan Fouant
USA

EMail: stefan.fouant@copperriverit.com

Robert Moskowitz
HTT Consulting
Oak Park, MI 48237
USA

EMail: rgm@labs.htt-consult.com

Nik Teague
Verisign
12061 Bluemont Way
Reston, VA 20190

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