

DOTS
Internet-Draft
Intended status: Informational
Expires: May 17, 2018

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

R. Moskowitz
HTT Consulting
N. Teague
Verisign
L. Xia
Huawei
K. Nishizuka
NTT Communications
November 13, 2017

Use cases for DDoS Open Threat Signaling
draft-ietf-dots-use-cases-09

Abstract

The DDoS Open Threat Signaling (DOTS) effort is intended to provide a protocol to facilitate interoperability between multivendor 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.

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 May 17, 2018.

Internet-Draft

DOTS Use Cases

November 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	4
3.	Use Cases Scenarios	4
3.1.	Inter-domain Use Cases	4
3.1.1.	Upstream Transit Providers Offering DDoS Mitigation Services	5
3.1.2.	Overlay MSSPs Offering DDoS Mitigation Services	10
3.1.3.	Examples of DOTS clients integrated with applications, services, and network infrastructure devices	12
3.2.	Intra-domain Use Cases	13
3.2.1.	Suppression of outbound DDoS traffic originating from a consumer broadband access network	14
3.2.2.	Home Network DDoS Detection Collaboration for ISP network management	16
3.2.3.	DDoS Orchestration	19
4.	Security Considerations	21
5.	IANA Considerations	22
6.	Acknowledgments	22
7.	References	22
7.1.	Normative References	22
7.2.	Informative References	22
	Authors' Addresses	22

[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 vendors, the primary goal of DOTS is to provide high-level interaction amongst differeing DDoS solutions, such as initiating or terminating DDoS mitigation assistance.

It should be noted that DOTS is not intended toperform orchestration functions; rather, DOTS is intended to allowdevices, 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.

[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](#)].

[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

Inter-domain DOTS deployment scenarios encompass two or more distinct spans of administrative control. A typical inter-domain DOTS deployment may consist of an endpoint network such as an Internet-connected enterprise requesting DDoS mitigation assistance from one or more upstream transit providers offering DDoS mitigation services, or from a topologically-distant MSSP offering cloud-based overlay DDoS mitigation services. DOTS may also be used to facilitate coordination of DDoS mitigation activities between mitigation providers.

Coordination between organizations making use of DOTS in such scenarios is necessary. Along with DOTS-specific tasks such as DOTS peering and validating the exchange of DOTS messaging between the relevant organizations, externalities relating to routing advertisements, authoritative DNS records (for DNS-based diversion), network access policies for DOTS nodes, service-level agreements (SLAs), and DDoS mitigation provisioning are required.

[3.1.1.](#) Upstream Transit Providers Offering DDoS Mitigation Services

The use-cases in this section reflect scenarios in which the immediate upstream transit network(s) offer DDoS mitigation services to their end-customers. The DOTS communications models in these use-cases are largely identical with that described in [Section 3.1.1.1](#), with the exceptions of [Section 3.1.1.4](#) and [Section 3.1.1.5](#). These variations are mainly of interest from the operational point of view, as they illustrate processes, procedures, and topological details which are externalities from the standpoint of the DOTS ecosystem, but which are of great importance for network operators and DDoS mitigation service providers.

[3.1.1.1.](#) End-customer with a single upstream transit provider offering DDoS mitigation services

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 requests 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 which is integrated with their DDoS mitigation orchestration system.

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.

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 transit

provider; 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.

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 organization which has deployed a DOTS-client-capable IDMS.

(b) The IDMS detects, classifies, and begins mitigating the DDoS attack.

(c) The IDMS determines that its capacity and/or capability to mitigate the DDoS attack is insufficient, and utilize its DOTS client functionality to send a DOTS mitigation service initiation request to one or more DOTS servers residing on the upstream transit networks. This DOTS mitigation service initiation request is automatically initiated by the DOTS client.

(d) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been configured to honor requests from the requesting DOTS client, and initiate situationally-appropriate DDoS mitigation service.

(e) The DOTS servers transmit a DOTS service status message to the DOTS client indicating that upstream DDoS mitigation service has been initiated.

(f) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the DOTS client.

(g) While DDoS mitigation services are active, the DOTS client may optionally regularly transmit DOTS mitigation efficacy updates to the DOTS server.

(h) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to the DOTS server.

(i) The DOTS server transmits a DOTS mitigation status update to the

DOTS client indicating that the DDoS attack has ceased.

(j) The DOTS client transmits a DOTS mitigation service termination request to the DOTS server.

(k) The DOTS server terminates DDoS mitigation service.

(l) The DOTS server transmits a DOTS mitigation status update to the DOTS client indicating that DDoS mitigation services have been terminated.

(m) The DOTS client transmits a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.1.2. End-customer with multiple upstream transit providers offering DDoS mitigation services

This scenario shares many characteristics with [Section 3.1.1.1](#), but with the key difference that the enterprise in question is multi-homed, i.e., has two or more upstream transit providers, and that they all provide DDoS mitigation services.

In most cases, the communications model for a multi-homed model would be the same as in the single-homed model, merely duplicated in parallel. However, if two or more of the upstream transit providers have entered into a mutual DDoS mitigation agreement and have established DOTS peering relationships, DDoS mitigation status messages may be exchanged between their respective DOTS servers in order to provide a more complete picture of the DDoS attack scope. This will also allow for either automated or operator-assisted programmatic cooperative DDoS mitigation activities on the part of the DOTS-peered transit providers.

The additional DOTS communications between the upstream transit providers will consist of mitigation start, mitigation status, and mitigation termination messages.

3.1.1.3. End-customer with multiple upstream transit providers, but

only a single upstream transit provider offering DDoS mitigation services

This scenario is similar to that outlined in [Section 3.1.1.2](#); however, only one of the upstream transit providers in question offers DDoS mitigation services. In this situation, the enterprise would cease advertising the relevant network prefixes via the transit providers which do not provide DDoS mitigation services - or, in the case where the enterprise does not control its own routing, request that the upstream transit providers which do not offer DDoS mitigation services stop advertising the relevant network prefixes on their behalf.

Once it has been determined that the DDoS attack has ceased, the enterprise once again announces the relevant routes to the upstream transit providers which do not offer DDoS mitigation services, or requests that they resume announcing the relevant routes on behalf of the enterprise.

Note that falling back to a single transit provider has the effect of reducing available inbound transit bandwidth during a DDoS attack. Without proper planning and sufficient provisioning of both the link capacity and DDoS mitigation capacity of the sole transit provider offering DDoS mitigation services, this reduction of available bandwidth could lead to network link congestion caused by legitimate inbound network traffic. Therefore, careful planning and provisioning of both upstream transit bandwidth as well as DDoS mitigation capacity is required in scenarios of this nature.

The withdrawal and announcement of routing prefixes described in this use-case falls outside the scope of DOTS, although they could conceivably be triggered as a result of provider-specific orchestration triggered by the receipt of specific DOTS messages from the enterprise in question.

[3.1.1.4](#). End-customer with an upstream transit provider offering DDoS mitigation services on an ongoing basis

This scenario is similar to that described in [Section 3.1.1.1](#), except that due to the lack of a capable local IDMS solution, strict uptime requirements, or other considerations, the end-customer has previously arranged for continuous active DDoS mitigation coverage; thus, the concepts of initiating or terminating DDoS mitigation services do not apply.

During an attack, the DOTS server on the upstream transit provider network immediately notifies the DOTS client on the enterprise

network about the attack and periodically signals the DOTS client in order to provide mitigation status information and related information. The end-customer may then use this information to add capacity or implement configuration changes intended to increase resilience, to plan maintenance windows or changes to routing policies, etc.

[3.1.1.5](#). End-customer with an upstream transit provider offering DDoS mitigation services, but DDoS mitigation service is refused

This scenario is similar to that described in [Section 3.1.1.1](#), except that the upstream transit provider is unable to honor the DDoS mitigation service request from the end-customer due to mitigation capacity limitations, technical faults, or other circumstances. The end-customer DOTS client may be configured to make repeated attempts to engage DDoS mitigation service from the refusing provider, or alternately may be configured to request DDoS mitigation service from alternate providers.

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 organization which has deployed a DOTS-client-capable IDMS.
- (b) The IDMS detects, classifies, and begins mitigating the DDoS attack.
- (c) The IDMS determines that its capacity and/or capability to mitigate the DDoS attack is insufficient, and utilizes its DOTS client functionality to send a DOTS mitigation service initiation request to a DOTS server residing on an upstream transit network.
- (d) The DOTS servers which receives the DOTS mitigation service initiation requests determines that it has been configured to honor requests from the requesting DOTS client, and attempts to initiate situationally-appropriate DDoS mitigation service.
- (e) The DDoS mitigators on the upstream network report back to the DOTS server that they are unable to initiate DDoS mitigation service for the requesting organization due to mitigation capacity constraints, bandwidth constraints, functionality constraints, hardware casualties, or other impediments.
- (f) The DOTS servers transmits a DOTS service status message to the DOTS client indicating that upstream DDoS mitigation service cannot

be initiated as requested.

(g) The DOTS client may optionally regularly re-transmit DOTS mitigation status request messages to the DOTS server until receiving acknowledgement that DDoS mitigation services have been initiated.

(h) The DOTS client may optionally transmit a DOTS mitigation service initiation request to a DOTS server associated with a configured fallback upstream DDoS mitigation service. Multiple fallback DDoS mitigation services may optionally be configured.

(i) The process describe above cyclically continues until the DDoS mitigation service request is fulfilled; the DOTS client determines that the DDoS attack volume has decreased to a level and/or complexity which it can successfully mitigate; the DDoS attack has ceased; or manual intervention by personnel of the requesting organization has taken place.

[3.1.2.](#) Overlay MSSPs Offering DDoS Mitigation Services

The use-cases in this section reflect scenarios in which topologically-distant managed security service providers (MSSPs) offer DDoS mitigation services to end-customers. The DOTS communications models in these use-cases are largely identical with that described in [Section 3.1.1.1](#). These variations are mainly of interest from the operational point of view, as they illustrate processes, procedures, and topological details which are externalities from the standpoint of the DOTS ecosystem, but which are of great importance for network operators and DDoS mitigation service providers.

[3.1.2.1.](#) End-customer with an overlay DDoS mitigation managed security service provider (MSSP)

This use case details an enterprise that has a local DDoS detection and classification capability and may or may not have an on-premise mitigation capability. The enterprise is contracted with an overlay DDoS mitigation MSSP, topologically distant from the enterprise network (i.e., not a direct upstream transit provider), which can redirect (divert) traffic away from the enterprise, provide DDoS mitigation services, and then forward (re-inject) legitimate

traffic to the enterprise on an on-demand basis. In this scenario, diversion of Internet traffic destined for the enterprise network into the overlay DDoS mitigation MSSP network is typically accomplished via eBGP announcements of the relevant enterprise network CIDR blocks, or via authoritative DNS subdomain-based mechanisms (other mechanisms are not precluded, these are merely the most common ones in use today).

The enterprise determines thresholds at which a request for mitigation is triggered indicating to the MSSP that inbound network traffic should be diverted into the MSSP network and that DDoS mitigation should be initiated. The enterprise may also elect to manually request diversion and mitigation via the MSSP network as desired.

The communications required to initiate, manage, and terminate active DDoS mitigation by the MSSP is performed using DOTS. The enterprise DDoS detection/classification system implements a DOTS client, while the MSSP implements a DOTS server integrated with its DDoS mitigation orchestration system. One or more out-of-band methods for initiating a mitigation request, such as a Web portal, a smartphone app, or voice support hotline, may also be made available by the MSSP.

When an attack is detected, an automated or manual DOTS mitigation request is generated by the enterprise and sent to the MSSP. The enterprise DOTS mitigation request is processed by the MSSP DOTS server, which validates the origin of the request and passes it to the MSSP DDoS mitigation orchestration system, which then initiates active DDoS mitigation. This action will usually involve the diversion of all network traffic destined for the targeted enterprise into the MSSP DDoS mitigation network, where it will be subjected to further scrutiny, with DDoS attack traffic filtered by the MSSP. Successful mitigation of the DDoS attack will not only result preserving the availability of services and applications resident on the enterprise network, but will also prevent DDoS attack traffic from ingressing the networks of the enterprise upstream transit providers and/or peers.

The MSSP should signal via DOTS to the enterprise that a mitigation request has been received and acted upon, and should also include an

update of the mitigation status. The MSSP may respond periodically with additional updates on the mitigation status to in order to enable the enterprise to make an informed decision on whether to maintain or terminate 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 also 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 DDoS monitoring and analysis service to customers whose networks may be protected by any one of a number of third-party providers. The enterprise in question may integrate with these third-party providers using DOTS and signal accordingly when a customer is attacked - the MSSP may then manage the life-cycle of the attack mitigation on behalf of the enterprise.

[3.1.2.2.](#) MSSP as an end-customer requesting overflow DDoS mitigation assistance from other MSSPs

This is a more complex use-case involving multiple DDoS MSSPs, whether transit operators, overlay MSSPs, or both. In this scenario, an MSSP has entered into a pre-arranged DDoS mitigation assistance agreement with one or more other DDoS MSSPs 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 DDoS MSSP to mitigate the attack on its own.

BGP-based diversion (including relevant Letters of Authorization, or LoAs), DNS-based diversion (including relevant LoAs), traffic re-injection mechanisms such as Generic Routing Encapsulation (GRE) tunnels, provisioning of DDoS orchestration systems, et. al,. must be arranged in advance between the DDoS MSSPs which are parties to the agreement. They should also be tested on a regular basis.

When a DDoS MSSP which is party to the agreement is nearing its capacity or ability to mitigate a given DDoS attack traffic, the DOTS client integrated with the MSSP DDoS mitigation orchestration system signals partner MSSPs to initiate network traffic diversion and DDoS mitigation activities. Ongoing attack and mitigation status messages may be passed between the DDoS MSSPs, and between the requesting MSSP

and the ultimate end-customer of the attack.

Once the requesting DDoS MSSP is confident that the DDoS attack has either ceased or has fallen to levels of traffic/complexity which they can handle on their own, the requesting DDoS MSSP DOTS client sends mitigation termination requests to the participating overflow DDoS MSSPs.

[3.1.3.](#) Examples of DOTS clients integrated with applications, services, and network infrastructure devices

The use-cases in this section reflect scenarios in which DOTS clients are integrated into devices and services which have heretofore been unable to request DDoS mitigation services on an as-needed basis. The DOTS communications models in these use-cases are largely identical with those described in [interdomain-use-cases](#). These variations are mainly of interest because they illustrate how DOTS allows DDoS mitigation services to expand in scope and utility.

[3.1.3.1.](#) End-customer operating an application or service with an integrated DOTS client

In this scenario, a Web server has a built-in mechanism to detect and classify DDoS attacks, which also incorporates a DOTS client. When an attack is detected, the self-defense mechanism is activated, and local DDoS mitigation is initiated.

The DOTS client built into the Web server has been configured to request DDoS mitigation services from an upstream transit provider or overlay MSSP once specific attack traffic thresholds have been reached, or certain network traffic conditions prevail. Once attack traffic subsides below configured thresholds and/or the specified network traffic conditions no longer apply, the DOTS client requests the termination of DDoS mitigation services.

[3.1.3.2.](#) End-customer operating a CPE network infrastructure device with an integrated DOTS client

Similar to the above use-case featuring applications or services with built-in DDoS attack detection/classification and DOTS client capabilities, in this scenario, an end-customer network infrastructure CPE device such as a router, layer-3 switch, firewall, IDS/IPS, Web application firewall or load-balancer incorporates both the functionality required to detect and classify incoming DDoS attacks as well as DOTS client functionality.

3.1.3.3. End-customer with an out-of-band smartphone application featuring DOTS client capabilities

This scenario would typically apply in a small office or home office (SOHO) setting, where the end-customer does not have CPE equipment or software capable of detecting/classifying/mitigating DDoS attack, yet still has a requirement for on-demand DDoS mitigation services. A smartphone application containing a DOTS client would be provided by the upstream transit mitigation provider or overlay DDoS MSSP to the SOHO end-customer; this application would allow a manual 'panic-button' to request the initiation and termination of DDoS mitigation services. The end-customer DOTS client application will display mitigation status information while DDoS mitigation activities are taking place.

3.2. Intra-domain Use Cases

While many of the DOTS-specific elements of inter-domain DOTS deployment scenarios apply to intra-domain scenarios, it is expected that many externalities such as coordination of and authorization for routing advertisements and authoritative DNS updates may be automated

to a higher degree than is practicable in inter-domain scenarios, given that the scope of required activities and authorizations are confined to a single organization. In theory, provisioning and change-control related both to DOTS itself as well as relevant externalities may require less administrative overhead and less implementation lead-times.

The scope of potential DDoS mitigation actions may also be broader in intra-organizational scenarios, as presumably an organization will have a higher degree of autonomy with regards to both technically and administratively feasible activities.

The DOTS communications model in these scenarios are largely identical to that described in [Section 3.1](#), except that all the DOTS communications take place within the same span of administrative control and the same network. These variations are mainly of interest from the operational point of view, as they illustrate processes, procedures, and topological details which are externalities from the standpoint of the DOTS ecosystem, but which are of great importance for network operators and DDoS mitigation service providers.

[3.2.1](#). Suppression of outbound DDoS traffic originating from a consumer broadband access network

While most DDoS defenses concentrate on inbound DDoS attacks ingressing from direct peering links or upstream transit providers, the DDoS attack traffic in question originates from one or more Internet-connected networks. In some cases, compromised devices residing on the local networks of broadband access customers are used to directly generate this DDoS attack traffic; in others, misconfigured devices residing on said local customer networks are exploited by attackers to launch reflection/amplification DDoS attacks. In either scenario, the outbound DDoS traffic emanating from these devices can be just as disruptive as an inbound DDoS attack, and can cause disruption for substantial proportions of the broadband access network operator's customer base.

Some broadband access network operators provide CPE devices (DSL modems/routers, cablemodems, FTTH routers, etc.) to their end-customers. Others allow end-customers to provide their own CPE devices. Many will either provide CPE devices or allow end-customers to supply their own.

Broadband access network operators typically have mechanisms to detect and classify both inbound and outbound DDoS attacks, utilizing flow telemetry exported from their peering/transit and customer aggregation routers. In the event of an outbound DDoS attack, they

may make use of internally-developed systems which leverage their subscriber-management systems to de-provision end-customers who are sourcing outbound DDoS traffic; in some cases, they may have implemented quarantine systems to block all outbound traffic sourced

from the offending end-customers. In either case, the perceived disruption of the end-customer's Internet access often prompts a help-desk call, which erodes the margins of the broadband access provider and can cause end-customer dissatisfaction.

Increasingly, CPE devices themselves are targeted by attackers who exploit security flaws in these devices in order to compromise them and subsume them into botnets, and then leverage them to launch outbound DDoS attacks. In all of the described scenarios, the end-customers are unaware that their computers and/or CPE devices have been compromised and are being used to launch outbound DDoS attacks - however, they may notice a degradation of their Internet connectivity as a result of outbound bandwidth consumption or other disruption.

By deploying DOTS-enabled telemetry systems and CPE devices (and possibly requiring DOTS functionality in customer-provided CPE devices), broadband access providers can utilize a standards-based mechanism to suppress outbound DDoS attack traffic while optionally allowing legitimate end-customer traffic to proceed unmolested.

In order to achieve this capability, the telemetry analysis system utilized by the broadband access provider must have DOTS client functionality, and the end-customer CPE devices must have DOTS server functionality. When the telemetry analysis system detects and classifies an outbound DDoS attack sourced from one or more end-customer networks/devices, the DOTS client of the telemetry analysis system sends a DOTS request to the DOTS server implemented on the CPE devices, requesting local mitigation assistance in suppressing either the identified outbound DDoS traffic, or all outbound traffic sourced from the end-customer networks/devices. The DOTS server residing within the CPE device(s) would then perform predefined actions such as implementing on-board access-control lists (ACLs) to suppress the outbound traffic in question and prevent it from leaving the local end-customer network(s).

Broadband access network operators may choose to implement a quarantine of all or selected network traffic originating from end-customer networks/devices which are sourcing outbound DDoS traffic, redirecting traffic from interactive applications such as Web browsers to an internal portal which informs the end-customer of the quarantine action, and providing instructions for self-remediation and/or helpdesk contact information.

Quarantine systems for broadband access networks are typically custom-developed and -maintained, and are generally deployed only by a relatively small number of broadband access providers with considerable internal software development and support capabilities. By requiring the manufacturers of operator-supplied CPE devices to implement DOTS server functionality, and requiring customer-provided CPE devices to feature DOTS server functionality, broadband access network operators who previously could not afford the development expense of creating custom quarantine systems to integrate DOTS-enabled network telemetry systems to act as DOTS clients and perform effective quarantine of end-customer networks and devices until such time as they have been remediated.

The DOTS communications model in this scenario resembles that described in [Section 3.1.1.1](#), except that all the DOTS communications take place within the same span of administrative control and the same network.

[3.2.2](#). Home Network 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

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.

[3.2.3.](#) 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 DDoS mitigation systems 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 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 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 ([Section 3.1.1.1](#)) or an MSSP ([Section 3.1.2.1](#)).

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.

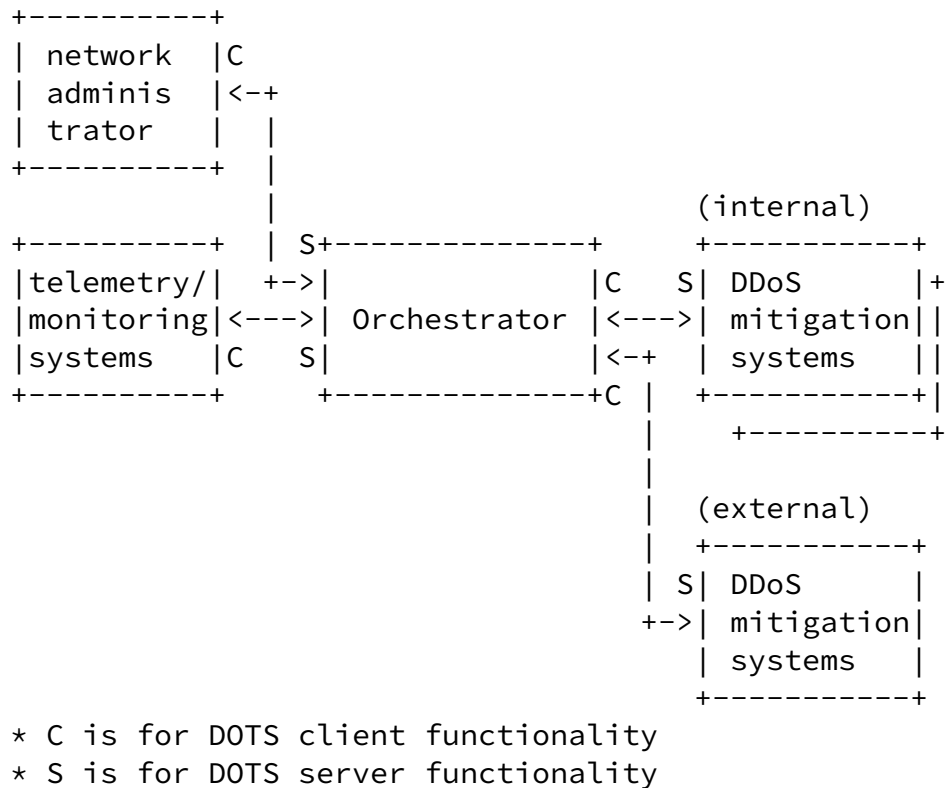


Figure 1: DDoS Orchestration

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 3.1.1.1](#) and [Section 3.1.2.1](#). 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 four primary forms of attack: DOTS agent impersonation, traffic injection, signal-blocking, and DDoS attacks. Associated security requirements and additional ones are defined in [[I-D.ietf-dots-requirements](#)].

These security risks can be managed through current secure communications best practices. DOTS is not subject to anything new or unique in this area.

[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; 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-07](#) (work in progress), October 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, <<https://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, <<https://www.rfc-editor.org/info/rfc7368>>.

Authors' Addresses

Roland Dobbins
Arbor Networks
Singapore

E-Mail: rdobbins@arbor.net

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

