

---

Stream: Internet Engineering Task Force (IETF)  
RFC: [8903](#)  
Category: Informational  
Published: May 2021  
ISSN: 2070-1721  
Authors:  
R. Dobbins    D. Migault    R. Moskowitz    N. Teague    L. Xia  
*Netscout, Inc.*    *Ericsson*    *HTT Consulting*    *Iron Mountain Data Centers*    *Huawei*  
K. Nishizuka  
*NTT Communications*

# RFC 8903

## Use Cases for DDoS Open Threat Signaling

---

### Abstract

The DDoS Open Threat Signaling (DOTS) effort is intended to provide protocols to facilitate interoperability across disparate DDoS Mitigation solutions. This document presents sample use cases that describe the interactions expected between the DOTS components as well as DOTS messaging exchanges. These use cases are meant to identify the interacting DOTS components, how they collaborate, and what the typical information to be exchanged is.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8903>.

### Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions

with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

- 1. [Introduction](#)
- 2. [Terminology and Acronyms](#)
- 3. [Use Cases](#)
  - 3.1. [Upstream DDoS Mitigation by an Upstream Internet Transit Provider](#)
  - 3.2. [DDoS Mitigation by a Third-Party DDoS Mitigation Service Provider](#)
  - 3.3. [DDoS Orchestration](#)
- 4. [Security Considerations](#)
- 5. [IANA Considerations](#)
- 6. [Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

## 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, coordinating 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 that are generally highly dependent on a timely reaction by the system.

The DDoS Open Threat Signaling (DOTS) effort is intended to specify protocols that facilitate interoperability between diverse DDoS Mitigation solutions and ensure greater integration in terms of attack detection, 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 detecting DDoS attacks, initiating/terminating DDoS Mitigation assistance, or requesting the status of a DDoS Mitigation.

This document provides sample use cases that provided input for the requirements [RFC8612] and design of the DOTS protocols [RFC8782][RFC8783]. The use cases are not exhaustive, and future use cases are expected to emerge as DOTS is adopted and evolves.

## 2. Terminology and Acronyms

This document makes use of the same terminology and definitions as [RFC8612]. In addition, it uses the terms defined below:

### DDoS Mitigation System (DMS):

A system that performs DDoS Mitigation. The DDoS Mitigation System may be composed of a cluster of hardware and/or software resources but could also involve an orchestrator that may make decisions, such as outsourcing some or all of the mitigation to another DDoS Mitigation System.

### DDoS Mitigation:

The action performed by the DDoS Mitigation System.

### DDoS Mitigation Service:

Designates a service provided to a customer to mitigate DDoS attacks. Each service subscription usually involve Service Level Agreement (SLA) that has to be met. It is the responsibility of the DDoS Service provider to instantiate the DDoS Mitigation System to meet these SLAs.

### DDoS Mitigation Service Provider:

Designates the administrative entity providing the DDoS Mitigation Service.

### Internet Transit Provider (ITP):

Designates the entity that delivers the traffic to a customer network. It can be an Internet Service Provider (ISP) or an upstream entity delivering the traffic to the ISP.

## 3. Use Cases

### 3.1. Upstream DDoS Mitigation by an Upstream Internet Transit Provider

This use case describes how an enterprise or a residential customer network may take advantage of a pre-existing relation with its ITP in order to mitigate a DDoS attack targeting its network.

For clarity of discussion, the targeted network is indicated as an enterprise network, but the same scenario applies to any downstream network, including residential and cloud hosting networks.

As the ITP provides connectivity to the enterprise network, it is already on the path of the inbound and outbound traffic of the enterprise network and is well aware of the networking parameters associated with the enterprise network WAN connectivity. This eases both the configuration and the instantiation of a DDoS Mitigation Service.

This section considers two kinds of DDoS Mitigation Service between an enterprise network and an ITP:

- The upstream ITP may instantiate a DMS upon receiving a request from the enterprise network. This typically corresponds to a 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 to the enterprise DMS to mitigate this at the source.

The two scenarios, though different, have similar interactions between the DOTS client and server. For the sake of simplicity, only the first scenario will be detailed in this section. Nevertheless, the second scenario is also in scope for DOTS.

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 Voice over IP (VoIP) servers has a DMS deployed to protect those servers and applications from DDoS attacks. In addition to on-premise DDoS defense capabilities, the enterprise has contracted with its ITP for DDoS Mitigation Services when attacks threaten to overwhelm the bandwidth of their WAN link (s).

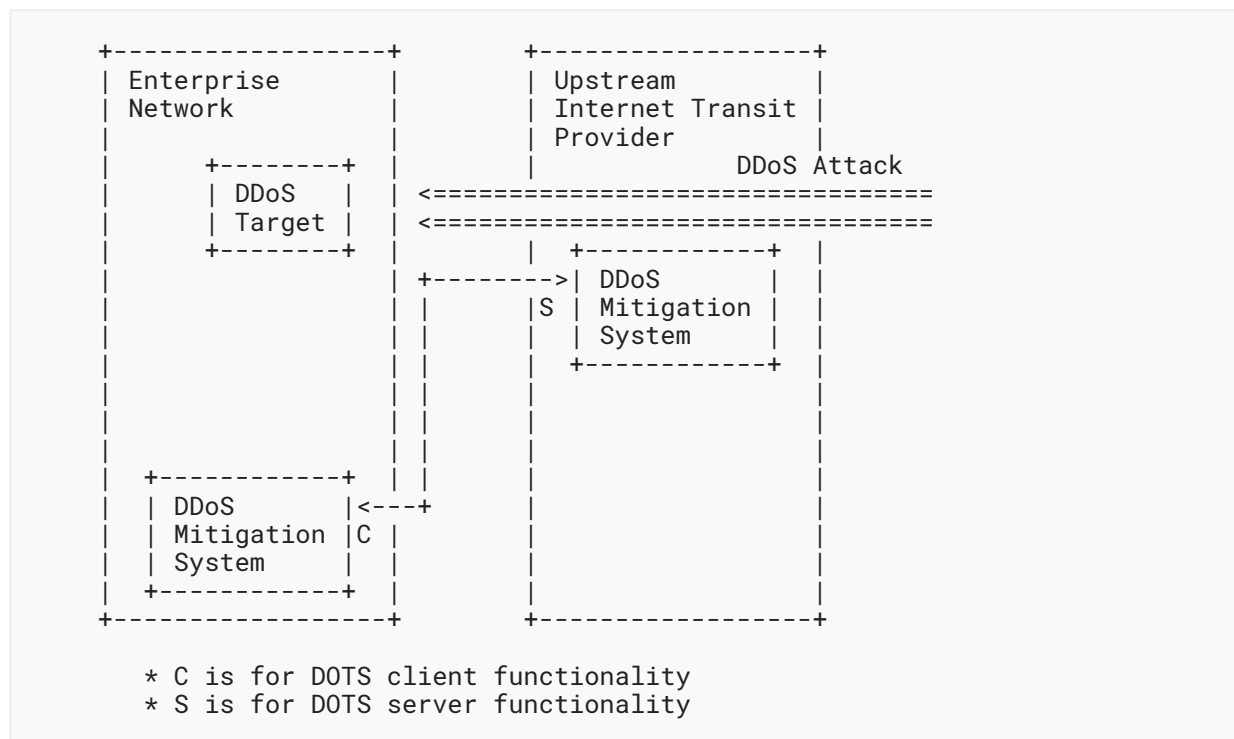


Figure 1: Upstream Internet Transit Provider DDoS Mitigation

The enterprise DMS is configured such that if the incoming Internet traffic volume exceeds 50% of the provisioned upstream Internet WAN link capacity, the DMS will request DDoS Mitigation assistance from the upstream transit provider. More sophisticated detection means may be considered as well.

The requests to trigger, manage, and finalize a DDoS Mitigation between the enterprise DMS and the ITP are made using DOTS. The enterprise DMS implements a DOTS client while the ITP implements a DOTS server, which is integrated with their DMS in this example.

When the enterprise DMS locally detects an inbound DDoS attack targeting its resources (e.g., servers, hosts, or applications), it immediately begins a DDoS Mitigation.

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

Over the course of the attack, the DOTS server of the ITP periodically informs the DOTS client on the mitigation status, statistics related to DDoS attack traffic mitigation, and related information. Once the DDoS attack has ended or decreased to a certain level that the enterprise DMS might handle by itself, the DOTS server signals the enterprise DMS DOTS client that the attack has subsided.

The DOTS client on the enterprise DMS then requests that the ITP 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:

1. A DDoS attack is initiated against resources of a network organization (here, the enterprise), which has deployed a DOTS-capable DMS -- typically a DOTS client.
2. The enterprise DMS detects, classifies, and begins the DDoS Mitigation.
3. The enterprise DMS determines that its capacity and/or capability to mitigate the DDoS attack is insufficient and sends a DOTS DDoS Mitigation request via its DOTS client to one or more DOTS servers residing on the upstream ITP.
4. The DOTS server, which receives the DOTS Mitigation request, determines that it has been configured to honor requests from the requesting DOTS client and does so by orchestrating its own DMS.
5. While the DDoS Mitigation is active, the DOTS server regularly transmits DOTS DDoS Mitigation status updates to the DOTS client.
6. Informed by the DOTS server status update that the attack has ended or subsided, the DOTS client transmits a DOTS DDoS Mitigation termination request to the DOTS server.
7. The DOTS server terminates DDoS Mitigation and sends the notification to the DOTS client.

Note that communications between the enterprise DOTS client and the upstream ITP DOTS server may take place in band within the main Internet WAN 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 third-party wireless 4G network connectivity.

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

Note also that when the enterprise is multihomed and connected to multiple upstream ITPs, each ITP is only able to provide a DDoS Mitigation Service for the traffic it transits. As a result, the enterprise network may be required to coordinate the various DDoS Mitigation Services associated with each link. More multihoming considerations are discussed in [\[DOTS-MULTIHOMING\]](#).

### **3.2. DDoS Mitigation by a Third-Party DDoS Mitigation Service Provider**

This use case differs from the previous use case described in [Section 3.1](#) in that the DDoS Mitigation Service is not provided by an upstream ITP. In other words, as represented in [Figure 2](#), the traffic is not forwarded through the DDoS Mitigation Service Provider by default. In order to steer the traffic to the DDoS Mitigation Service Provider, some network configuration changes are required. As such, this use case is likely to apply to large enterprises or large data centers but, as for the other use cases, is not exclusively limited to them.

Another typical scenario for this use case is for there to be a relationship between DDoS Mitigation Service Providers, forming an overlay of DMS. When a DDoS Mitigation Service Provider mitigating a DDoS attack reaches its resource capacity, it may choose to delegate the DDoS Mitigation to another DDoS Mitigation Service Provider.

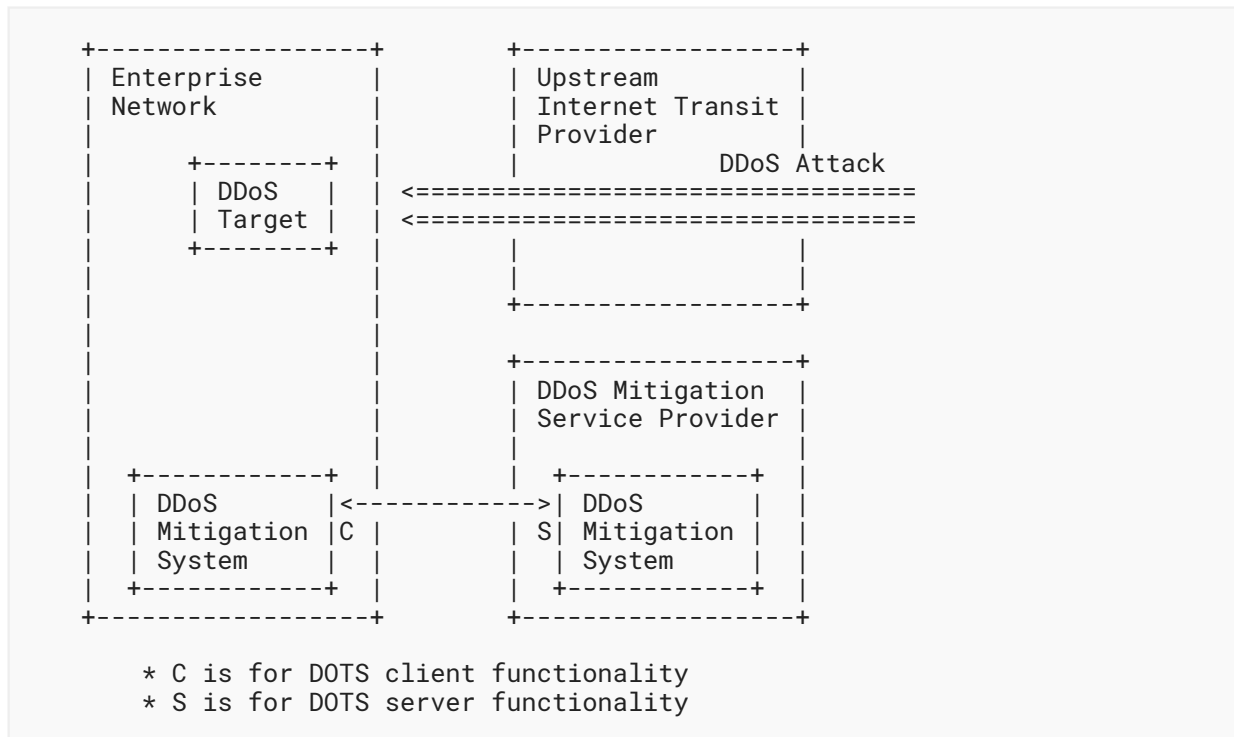


Figure 2: DDoS Mitigation between an Enterprise Network and a Third-Party DDoS Mitigation Service Provider

In this scenario, an enterprise network has entered into a prearranged DDoS Mitigation assistance agreement with one or more third-party 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 the enterprise or any other given DMS to mitigate the attack on its own.

The prearrangement typically includes agreement on the mechanisms used to redirect the traffic to the DDoS Mitigation Service Provider, as well as the mechanism to re-inject the traffic back to the Enterprise Network. Redirection to the DDoS Mitigation Service Provider typically involves BGP prefix announcement or DNS redirection, while re-injection of the scrubbed traffic to the enterprise network may be performed via tunneling mechanisms (e.g., GRE). The exact mechanisms used for traffic steering are out of scope of DOTS but will need to be prearranged, while in some contexts such changes could be detected and considered as an attack.

In some cases, the communication between the enterprise DOTS client and the DOTS server of the DDoS Mitigation Service Provider may go through the ITP carrying the DDoS attack, which would affect the communication. On the other hand, the communication between the DOTS client and DOTS server may take a path that is not undergoing a DDoS attack.

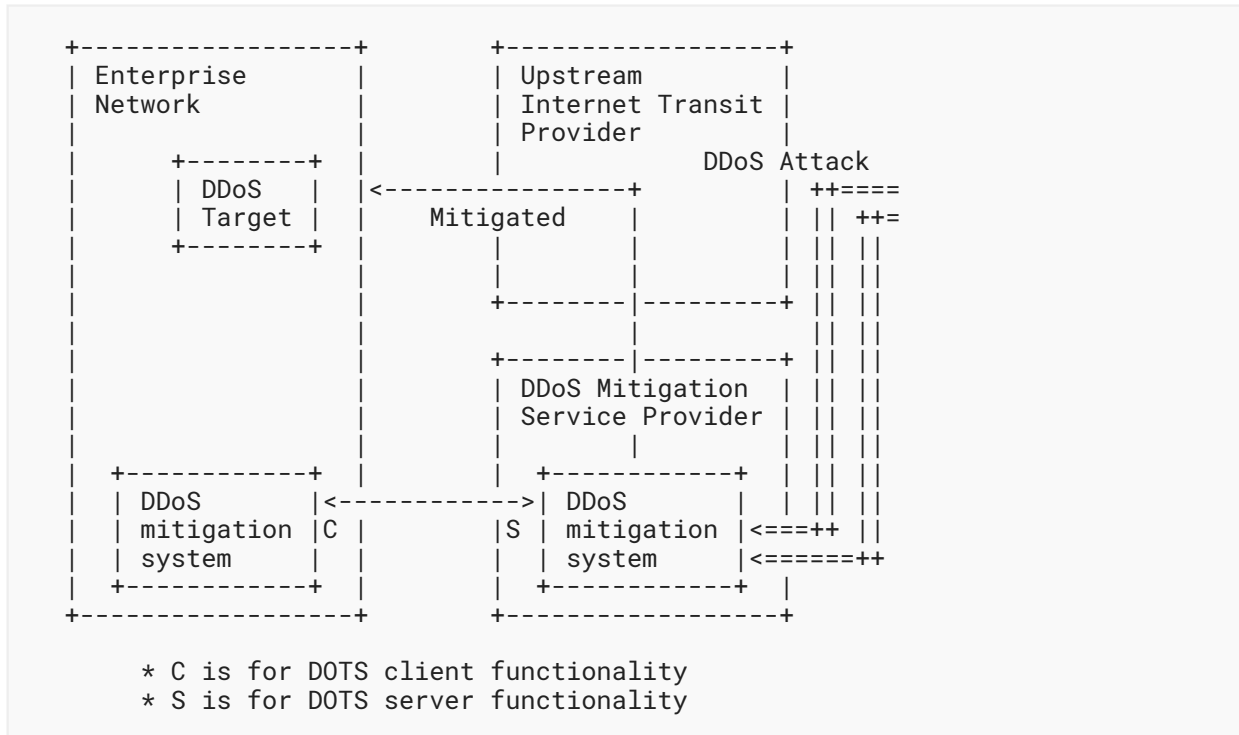


Figure 3: Redirection to a DDoS Mitigation Service Provider

When the enterprise network is under attack or at least is reaching its capacity or ability to mitigate a given DDoS attack, 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 using DOTS. If the DDoS attack has stopped or the severity of the attack has subsided, the DOTS client can request that the DDoS Mitigation Service Provider terminate the DDoS Mitigation.

### 3.3. DDoS Orchestration

In this use case, one or more DDoS telemetry systems or monitoring devices monitor a network -- typically an ISP network, an enterprise network, or a data center. Upon detection of a DDoS attack, these DDoS telemetry systems alert an orchestrator in charge of coordinating the various DMSs within the domain. The DDoS telemetry systems may be configured to provide required information, such as a preliminary analysis of the observation, to the orchestrator.

The orchestrator analyzes the various sets of information it receives from DDoS telemetry systems and initiates one or more DDoS Mitigation strategies. For example, the orchestrator could select the DMS in the enterprise network or one provided by the ITP.

DMS selection and DDoS Mitigation techniques may depend on the type of the DDoS attack. In some cases, a manual confirmation or selection may also be required to choose a proposed strategy to initiate a DDoS Mitigation. The DDoS Mitigation may consist of multiple steps such as



configuring the network or updating already-instantiated DDoS Mitigation functions. Eventually, the coordination of the mitigation may involve external DDoS Mitigation resources such as a transit provider or a third-party DDoS Mitigation Service Provider.

The communication used to trigger a DDoS Mitigation between the DDoS telemetry and monitoring systems and the orchestrator is performed using DOTS. The DDoS telemetry system 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 uses, for example, a web interface that interacts with a DOTS client, while the orchestrator implements a DOTS server.

The communication between the orchestrator and the DMSs is performed using DOTS. The orchestrator implements a DOTS client while the DMSs implement a DOTS server.

The configuration aspects of each DMS, as well as the instantiations of DDoS Mitigation functions or network configuration, are not part of DOTS. Similarly, the discovery of available DDoS Mitigation functions is not part of DOTS and, as such, is out of scope.

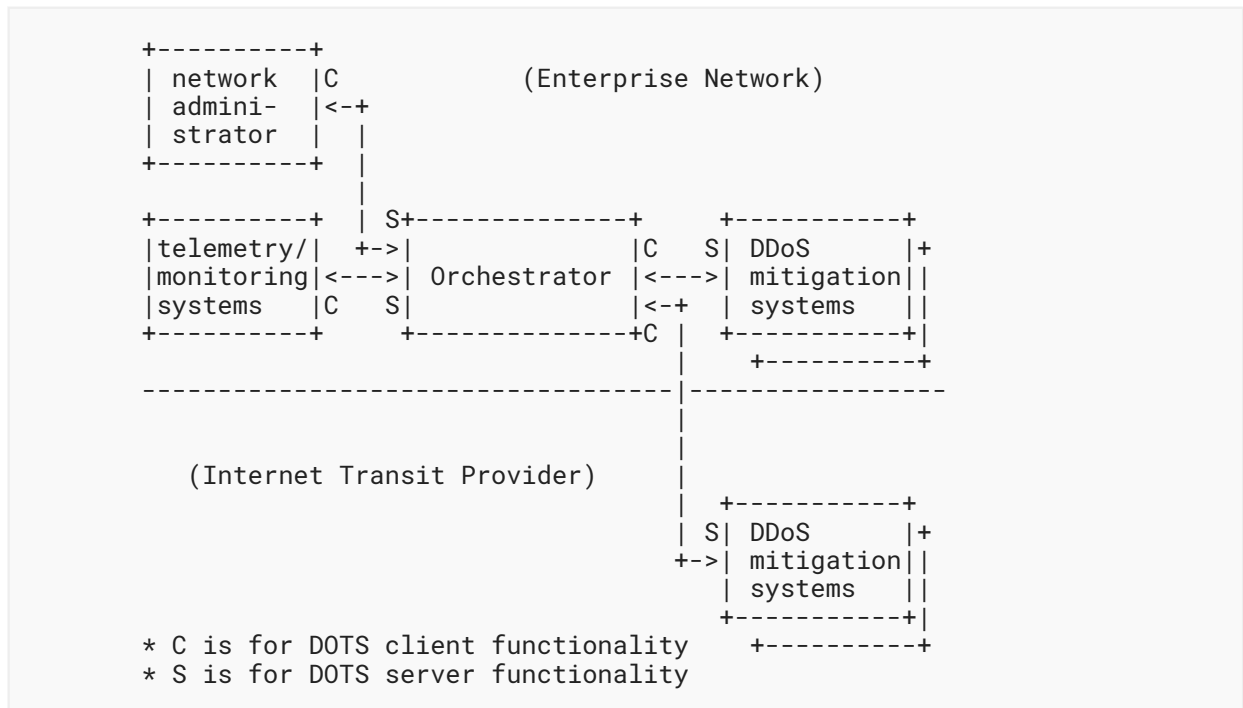


Figure 4: DDoS Orchestration

The DDoS telemetry systems monitor various aspects of the network traffic and perform some measurement tasks.

These systems are configured so that when an event or some measurement indicators reach a predefined level, their associated DOTS client sends a DOTS mitigation request to the orchestrator DOTS server. The DOTS mitigation request may be associated with some optional

mitigation hints to let the orchestrator know what has triggered the request. In particular, it is possible for something that looks like an attack locally to one telemetry system is not actually an attack when seen from the broader scope (e.g., of the orchestrator).

Upon receipt of the DOTS mitigation request from the DDoS telemetry system, the orchestrator DOTS server responds with an acknowledgment to avoid retransmission of the request for mitigation. The orchestrator may begin collecting additional fine-grained and specific information from various DDoS telemetry systems in order to correlate the measurements and provide an analysis of the event. Eventually, the orchestrator may ask for additional information from the DDoS telemetry system; however, the collection of this information is out of scope of 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, for example, a web interface. If the network administrator decides to start the mitigation, the network administrator triggers the DDoS Mitigation request using, for example, a web interface of a DOTS client communicating to the orchestrator DOTS server. This request is expected to be associated with a context that provides sufficient information to the orchestrator DOTS server to infer, elaborate, and coordinate the appropriate DDoS Mitigation.

Upon receiving a request to mitigate a DDoS attack aimed at a target, the orchestrator may evaluate the volume of the attack as well as the value that the target represents. The orchestrator may select the DDoS Mitigation Service Provider based on the attack severity. It may also coordinate the DDoS Mitigation performed by the DDoS Mitigation Service Provider with some other tasks such as, for example, moving the target to another network so new sessions will not be impacted. The orchestrator requests a DDoS Mitigation by the selected DMSs via its DOTS client, as described in [Section 3.1](#).

The orchestrator DOTS client is notified that the DDoS Mitigation is effective by the selected DMSs. The orchestrator DOTS server returns this information to the network administrator.

Similarly, when the DDoS attack has stopped, the orchestrator DOTS client is notified and the orchestrator's DOTS server indicates the end of the DDoS Mitigation to the DDoS telemetry systems as well as to the network administrator.

In addition to the DDoS orchestration shown in [Figure 4](#), the selected DMS can return a mitigation request to the orchestrator as an offloading. For example, when the DDoS attack becomes severe and the DMS's utilization rate reaches its maximum capacity, the DMS can send mitigation requests with additional hints, such as its blocked traffic information, to the orchestrator. Then the orchestrator can take further actions such as requesting forwarding nodes (e.g., routers) to filter the traffic. In this case, the DMS implements a DOTS client while the orchestrator implements a DOTS server. Similar to other DOTS use cases, the offloading scenario assumes that some validation checks are followed by the DMS, the orchestrator, or both (e.g., avoid exhausting the resources of the forwarding nodes or inadvertent disruption of legitimate services). These validation checks are part of the mitigation and are therefore out of the scope of the document.

## 4. Security Considerations

The document does not describe any protocol, though there are still a few high-level security considerations to discuss.

DOTS is at risk from three primary attacks: DOTS agent impersonation, traffic injection, and signaling blocking.

Impersonation and traffic injection mitigation can be mitigated through current secure communications best practices, including mutual authentication. Preconfigured mitigation steps to take on the loss of keepalive traffic can partially mitigate signal blocking. But in general, it is impossible to comprehensively defend against an attacker that can selectively block any or all traffic. Alternate communication paths that are (hopefully) not subject to blocking by the attacker in question is another potential mitigation.

Additional details of DOTS security requirements can be found in [RFC8612].

Service disruption may be experienced if inadequate mitigation actions are applied. These considerations are out of the scope of DOTS.

## 5. IANA Considerations

This document has no IANA actions.

## 6. Informative References

- [DOTS-MULTIHOMING] Boucadair, M., Reddy, T., and W. Pan, "Multi-homing Deployment Considerations for Distributed-Denial-of-Service Open Threat Signaling (DOTS)", Work in Progress, Internet-Draft, draft-ietf-dots-multihoming-06, 25 May 2021, <<https://tools.ietf.org/html/draft-ietf-dots-multihoming-06>>.
- [RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <<https://www.rfc-editor.org/info/rfc8612>>.
- [RFC8782] Reddy, K. T., Ed., Boucadair, M., Ed., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", RFC 8782, DOI 10.17487/RFC8782, May 2020, <<https://www.rfc-editor.org/info/rfc8782>>.
- [RFC8783] Boucadair, M., Ed. and T. Reddy, K., Ed., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", RFC 8783, DOI 10.17487/RFC8783, May 2020, <<https://www.rfc-editor.org/info/rfc8783>>.

## Acknowledgments

The authors would like to thank, among others, Tirumaleswar Reddy.K, Andrew Mortensen, Mohamed Boucadair, Artyom Gavrichenkov, Jon Shallow, Yuuhei Hayashi, Elwyn Davies, the DOTS WG Chairs (at the time of writing) Roman Danyliw and Tobias Gondrom, as well as the Security AD Benjamin Kaduk for their valuable feedback.

We also would like to thank Stephan Fouant, who was one of the initial coauthors of the documents.

## Authors' Addresses

### **Roland Dobbins**

Netscout, Inc.  
Singapore  
Email: [roland.dobbins@netscout.com](mailto:roland.dobbins@netscout.com)

### **Daniel Migault**

Ericsson  
8275 Trans Canada Route  
Saint Laurent, Quebec 4S 0B6  
Canada  
Email: [daniel.migault@ericsson.com](mailto:daniel.migault@ericsson.com)

### **Robert Moskowitz**

HTT Consulting  
Oak Park, MI 48237  
United States of America  
Email: [rgm@labs.htt-consult.com](mailto:rgm@labs.htt-consult.com)

### **Nik Teague**

Iron Mountain Data Centers  
United Kingdom  
Email: [nteague@ironmountain.co.uk](mailto:nteague@ironmountain.co.uk)

### **Liang Xia**

Huawei  
No. 101, Software Avenue, Yuhuatai District  
Nanjing  
China  
Email: [Frank.xialiang@huawei.com](mailto:Frank.xialiang@huawei.com)

**Kaname Nishizuka**

NTT Communications

GranPark 16F

3-4-1 Shibaura, Minato-ku, Tokyo

108-8118

Japan

Email: [kaname@nttv6.jp](mailto:kaname@nttv6.jp)