DOTS WG
Internet-Draft
Intended status: Informational
Expires: September 22, 2016

R. Dobbins, Ed.
Arbor Networks
S. Fouant
Corero Network Security
D. Migault
Ericsson
R. Moskowitz
HTT Consulting
N. Teague
Verisign Inc
L. Xia
Huawei

March 21, 2016

Use cases for DDoS Open Threat Signaling
draft-ietf-dots-use-cases-01.txt

Abstract

This document delineates principal and ancillary use cases for DDoS Open Threat Signaling (DOTS), a communications protocol intended to facilitate the programmatic, coordinated mitigation of Distributed Denial of Service (DDoS) attacks via a standards-based mechanism. DOTS is purposely designed to support requests for DDoS mitigation services and status updates across inter-organizational administrative boundaries.

Table of Contents

1. Introduction

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

The DDoS Open Threat Signaling (DOTS) effort is intended to facilitate interoperability between DDoS solutions/services by providing a standards-based, programmatic communications mechanism for the invitation and termination of heterogeneous DDoS attack mitigation systems and services. This allows for a much higher degree of automation and concomitant efficacy and rapidity of DDoS attack mitigation involving multiple DDoS mitigation systems and services than is currently the norm, as well as providing additional benefits such as automatic DDoS mitigation service registration and provisioning. It should be noted that DOTS is not in and of itself intended to perform orchestration functions duplicative of the functionality being developed by the [I2NSF] WG; rather, DOTS is intended to allow devices, services, and applications to request mitigation assistance and receive mitigation status updates from systems of this nature.

This document provides an overview of common DDoS mitigation system/service deployment and operational models which are in use today, but which are currently limited in scope to a single vendor or service provider and are often highly manual in nature, which can lead to miscommunications, misconfigurations, and delays in bringing mitigation services to bear against an attack. The introduction of DOTS into these scenarios will reduce reaction times and the risks associated with manual processes, simplify the use of multiple types of DDoS mitigation systems and services as required, and make practical the simultaneous use multiple DDoS mitigation systems and services as circumstances warrant.

2. Terminology and Acronyms

2.1. Requirements Terminology

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

2.2. Acronyms

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

3. Use Cases

This section provides a high-level overview of likely use cases and deployment scenarios for DOTS-enabled DDoS mitigation services. It should be noted that DOTS servers may be standalone entities which, upon receiving a DOTS mitigation service request from a DOTS client, proceed to initiate DDoS mitigation service by communicating directly or indirectly with DDoS mitigators, and likewise terminate the service upon receipt of a DOTS service termination request; conversely, the DDoS mitigators themselves may incorporate DOTS servers and/or DOTS clients. The mechanisms by which DOTS servers initiate and terminate DDoS mitigation service with DDoS mitigators is beyond the scope of this document.

All of the primary use cases described in this section are derived from current, real-world DDoS mitigation functionality, capabilities, and operational models.

The posited ancillary use cases described in this section are reasonable and highly desirable extrapolations of the functionality of baseline DOTS capabilities, and are readily attainable in the near term.

Each of the primary and ancillary use cases described in this section may be read as involving one or more DDoS mitigation service providers; DOTS makes multi-provider coordinated DDoS defenses much more effective and practical due to abstraction of the particulars of a given DDoS mitigation service/solution set.

Both the primary and ancillary use cases may be facilitated by direct DOTS client - DOTS server communications or via DOTS relays deployed in order to aggregate DOTS mitigation service requests/responses, to mediate between stateless and stateful underlying transport protocols, to aggregate multiple DOTS requests and/or responses, to

filter DOTS requests and/or responses via configured policy mechanisms, or some combination of these functions.

All DOTS messages exchanged between the DOTS clients and DOTS servers in these use cases may be communicated directly between DOTS clients and servers, or mediated by one or more DOTS relays residing on the network of the originating network, the network where upstream DDoS mitigation service takes place, an intervening network or networks, or some combination of the above.

DOTS is intended to apply to both inter- and intra-domain DDoS attack mitigation scenarios. The technical and operational requirements for inter- and intra-domain DOTS communications are identical. The main difference is administrative in nature; although it should be noted that provisioning challenges which are typically associated with inter- domain DOTS communications relationships may also apply in intra- domain deployment scenarios, based upon organizational factors. All of the same complexities surrounding authentication and authorization can apply in both contexts, including considerations such as network access policies to allow DOTS communications, DOTS transport selection (including considerations of the implications of link congestion if a stateful DOTS transport option is selected), etc. Registration of well-known ports for DOTS transports per [RFC6335] should be considered in light of these challenges.

It should also be noted that DOTS does not directly ameliorate the various administrative challenges required for successful DDoS attack mitigation. Letters of authorization, RADB updates, DNS zone delegations, alteration of network access policies, technical configurations required to facilitate network traffic diversion and re-injection, etc., are all outside the scope of DOTS. DOTS may, however, prove useful in automating the registration of DOTS clients with DOTS servers, as well as in the automatic provisioning of situationally- appropriate DDoS defenses and countermeasures. This ancillary DOTS functionality is described in Section 3.2.

Many of the 'external' administrative challenges associated with establishing workable DDoS attack mitigation service may be addressed by work currently in progress in the I2RS and I2NSF WGs. Interested parties may wish to consider tracking those efforts, and coordination with both I2RS and I2NSF is highly desirable.

Note that all the use-cases in this document are universal in nature. They apply equally to endpoint networks, transit backbone providers, cloud providers, broadband access providers, ASPs, CDNs, etc. They are not specific to particular business models, topological models, or application types, and are deliberately generalizable. Both networks targeted for attack as well as any adjacent or topologically

distant networks involved in a given scenario may be either single- or multi-homed. In the accompanying vector illustrations incorporated into draft-ietf-dots-use-cases-01.pdf, specific business and topological models are described in order to provide context.

Likewise, both DOTS itself and the use cases described in this document are completely independent of technologies utilized for the detection, classification, traceback, and mitigation of DDoS attacks. Flow telemetry such as NetFlow and IPFIX, direct full-packet analysis, log-file analysis, indirection manual observation, etc. can and will be enablers for detection, classification and traceback. Intelligent DDoS mitigation systems (IDMSes), flowspec, S/RTBH, ACLs, and other network traffic manipulation tools and techniques may be used for DDoS attack mitigation. BGP, flowspec, DNS, inline deployment, and various 'NFV' technologies may be used for network traffic diversion into mitigation centers or devices in applicable scenarios; GRE, MPLS, 'NFV', inline deployment and other techniques may be utilized for 'cleaned' traffic re-injection to its intended destination.

The scope, format, and content of all DOTS message types cited in this document must be codified by the DOTS WG.

The following use cases are intended to inform the DOTS requirements described in [I-D.ietf-dots-requirements].

3.1. Primary Use Cases

3.1.1. Automatic or Operator-Assisted CPE or PE Mitigators Request Upstream DDoS Mitigation Services

One or more CPE or PE mitigators with DOTS client capabilities may be configured to signal to one or more DOTS servers in order to request upstream DDoS mitigation service initiation during an attack when DDoS attack volumes and/or attack characteristics exceed the capabilities of such CPE mitigators. DDoS mitigation service may be terminated either automatically or manually via a DOTS mitigation service termination request initiated by the mitigator when it has been determined that the DDoS attack has ended.

(a) A DDoS attack is initiated against online properties of an organization which has deployed DOTS-client-capable DDoS mitigators.

(b) CPE or PE DDoS mitigators detect, classify, and begin mitigating the DDoS attack.

(c) CPE or PE DDoS mitigators determine that their capacity and/or capability to mitigate the DDoS attack is insufficient, and utilize their DOTS client functionality to send a DOTS mitigation service initiation request to one or more DOTS servers residing on one or more upstream transit networks, peer networks, or overlay MSSP networks. This DOTS mitigation service initiation request may be automatically initiated by the CPE or PE DDoS mitigators, or may be manually triggered by personnel of the requesting organization in response to an alert from the mitigators (the mechanism by which this process takes place is beyond the scope of this document).

(d) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been configured to honor requests from the requesting CPE or PE mitigators, and initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(e) The DOTS servers transmit a DOTS service status message to the requesting CPE or PE mitigators 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 requesting CPE or PE mitigators.

(g) While DDoS mitigation services are active, the CPE or PE mitigators may optionally regularly transmit DOTS mitigation efficacy updates to the relevant DOTS servers.

(h) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to their respective DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(i) The DOTS servers transmit a DOTS mitigation status update to the CPE or PE mitigators indicating that the DDoS attack has ceased.

(j) The CPE or PE DDoS mitigators transmit a DOTS mitigation service termination request to the DOTS servers. This DOTS mitigation service termination request may be automatically initiated by the CPE or PE DDoS mitigators, or may be manually triggered by personnel of the requesting organization in response to an alert from the mitigators or a management system which monitors them (the mechanism by which this process takes place is beyond the scope of this document).

(k) The DOTS servers terminate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(l) The DOTS servers transmit a DOTS mitigation status update to the CPE or PE mitigators indicating that DDoS mitigation services have been terminated.

(m) The CPE or PE DDoS mitigators transmit a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.2. Automatic or Operator-Assisted CPE or PE Network Infrastructure Element Request to Upstream Mitigator

CPE or PE network infrastructure elements such as routers, switches, load-balancers, firewalls, 'IPSes', etc. which have the capability to detect and classify DDoS attacks and which have DOTS client capabilities may be configured to signal to one or more DOTS servers in order to request upstream DDoS mitigation service initiation during an attack. DDoS mitigation service may be terminated either automatically or manually via a DOTS mitigation service termination request initiated by the network element when it has been determined that the DDoS attack has ended.

In this use-case, the network elements involved are not engaged in mitigating DDoS attack traffic. They are signaling for upstream attack mitigation assistance. This can be an inter- or intra- domain use-case.

(a) A DDoS attack is initiated against online properties of an organization with DOTS-client-capable network infrastructure elements deployed.

(b) The network infrastructure elements utilize their DOTS client functionality to send a DOTS mitigation service initiation request to one or more DOTS servers residing on one or more upstream transit networks, peer networks, or overlay MSSP networks, either directly or via intermediate DOTS relays residing upon the requesting organization's network, the upstream mitigation provider's network, or both. The scope, format, and content of these messages must be codified by the DOTS WG. This DOTS mitigation service initiation request may be automatically initiated by the network infrastructure elements, or may be manually triggered by personnel of the requesting organization in response to an alert from the network elements or a management system which monitors them (the mechanism by which this process takes place is beyond the scope of this document).

(c) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been configured to honor requests from the requesting network infrastructure elements, and initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(d) The DOTS servers transmit a DOTS service status message to the requesting network infrastructure elements indicating that upstream DDoS mitigation service has been initiated.

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the requesting requesting network infrastructure elements.

(f) While DDoS mitigation services are active, the network infrastructure elements may optionally regularly transmit DOTS mitigation efficacy updates to the relevant DOTS servers.

(g) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to their respective DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(h) The DOTS servers transmit a DOTS mitigation status update to the network infrastructure elements indicating that the DDoS attack has ceased.

(i) The network infrastructure elements transmit a DOTS mitigation service termination request to the DOTS servers. This DOTS mitigation service termination request may be automatically initiated by the network infrastructure elements, or may be manually triggered by personnel of the requesting organization in response to an alert from the mitigators (the mechanism by which this process takes place is beyond the scope of this document).

(j) The DOTS servers terminate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(k) The DOTS servers transmit a DOTS mitigation status update to the network infrastructure elements indicating that DDoS mitigation services have been terminated.

(l) The network infrastructure elements transmit a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.3. Automatic or Operator-Assisted CPE or PE Attack Telemetry Detection/Classification System Request to Upstream Mitigator

CPE or PE Attack Telemetry Detection/Classification Systems which have DOTS client capabilities may be configured so that upon detecting and classifying a DDoS attack, they signal one or more DOTS servers in order to request upstream DDoS mitigation service initiation. DDoS mitigation service may be terminated either automatically or manually via a DOTS mitigation service termination request initiated by the Attack Telemetry Detection/Classification System when it has been determined that the DDoS attack has ended.

In this use-case, the Attack Telemetry Detection/Classification does not possess any inherent capability to mitigate DDoS attack traffic, and is signaling for upstream mitigation assistance. This can be an inter- or intra-domain use-case.

(a) A DDoS attack is initiated against online properties of an organization with DOTS-client-capable CPE or PE Attack Telemetry Detection/Classification Systems deployed.

(b) The CPE or PE Attack Telemetry Detection/Classification Systems utilize their DOTS client functionality to send a DOTS mitigation service initiation request to one or more DOTS servers residing on one or more upstream transit networks, peer networks, or overlay MSSP networks, either directly or via intermediate DOTS relays residing upon the requesting organization's network, the upstream mitigation provider's network, or both. This DOTS mitigation service initiation request may be automatically initiated by the CPE or PE Attack Telemetry Detection/Classification Systems, or may be manually triggered by personnel of the requesting organization in response to an alert from the CPE or PE Attack Telemetry Detection/Classification Systems (the mechanism by which this process takes place is beyond the scope of this document).

(c) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been configured to honor requests from the requesting CPE or PE Attack Telemetry Detection/Classification Systems, and initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(d) The DOTS servers transmit a DOTS service status message to the requesting CPE or PE Attack Telemetry Detection/Classification Systems indicating that upstream DDoS mitigation service has been initiated.

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the requesting CPE or PE Attack Telemetry Detection/Classification Systems.

(f) While DDoS mitigation services are active, the CPE or PE Attack Telemetry Detection/Classification Systems may optionally regularly transmit DOTS mitigation efficacy updates to the relevant DOTS servers.

(g) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to their respective DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(h) The DOTS servers transmit a DOTS mitigation status update to the CPE or PE Attack Telemetry Detection/Classification Systems indicating that the DDoS attack has ceased.

(i) The CPE or PE Attack Telemetry Detection/Classification Systems transmit a DOTS mitigation service termination request to the DOTS servers. This DOTS mitigation service termination request may be automatically initiated by the CPE or PE Attack Telemetry Detection/Classification Systems, or may be manually triggered by personnel of the requesting organization in response to an alert from the CPE or PE Attack Telemetry Detection/ Classification Systems (the mechanism by which this process takes place is beyond the scope of this document).

(j) The DOTS servers terminate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(k) The DOTS servers transmit a DOTS mitigation status update to the CPE or PE Attack Telemetry Detection/Classification Systems indicating that DDoS mitigation services have been terminated.

(l) The CPE or PE Attack Telemetry Detection/Classification Systems transmit a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.4. Automatic or Operator-Assisted Targeted Service/ Application Request to Upstream Mitigator

A service or application which is the target of a DDoS attack and which has the capability to detect and classify DDoS attacks (i.e, Apache mod_security [APACHE], BIND RRL [RRL], etc.) as well as DOTS client functionality may be configured so that upon detecting and

classifying a DDoS attack, it signals one or more DOTS servers in order to request upstream DDoS mitigation service initiation. DDoS mitigation service may be terminated either automatically or manually via a DOTS mitigation service termination request initiated by the service/application when it has been determined that the DDoS attack has ended.

In this use-case, the service/application does not possess inherent DDoS attack mitigation capabilities, and is signaling for upstream mitigation assistance. This can be an inter- or intra-domain use-case.

(a) A DDoS attack is initiated against online properties of an organization which include DOTS-client-capable services or applications that are the specific target(s) of the attack.

(b) The targeted services or applications utilize their DOTS client functionality to send a DOTS mitigation service initiation request to one or more DOTS servers residing on the same network as the services or applications, one or more upstream transit networks, peer networks, or overlay MSSP networks, either directly or via intermediate DOTS relays residing upon the requesting organization's network, the upstream mitigation provider's network, or both. This DOTS mitigation service initiation request may be automatically initiated by the targeted services or applications, or may be manually triggered by personnel of the requesting organization in response to an alert from the targeted services or applications or a system which monitors them (the mechanism by which this process takes place is beyond the scope of this document).

(c) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been provisioned to honor requests from the requesting services or applications, and initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(d) The DOTS servers transmit a DOTS service status message to the services or applications indicating that upstream DDoS mitigation service has been initiated

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the requesting services or applications.

(f) While DDoS mitigation services are active, the requesting services or applications may optionally regularly transmit DOTS mitigation efficacy updates to the relevant DOTS servers.

(g) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to their respective DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(h) The DOTS servers transmit a DOTS mitigation status update to the requesting services or applications indicating that the DDoS attack has ceased.

(i) The targeted services or applications transmit a DOTS mitigation service termination request to the DOTS servers. This DOTS mitigation service termination request may be automatically initiated by the targeted services or applications, or may be manually triggered by personnel of the requesting organization in response to an alert from a system which monitors them (the mechanism by which this process takes place is beyond the scope of this document).

(j) The DOTS servers terminate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(k) The DOTS servers transmit a DOTS mitigation status update to the targeted services or applications indicating that DDoS mitigation services have been terminated.

(l) The targeted services or applications transmit a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.5. Manual Web Portal Request to Upstream Mitigator

A Web portal which has DOTS client capabilities has been configured in order to allow authorized personnel of organizations which are targeted by DDoS attacks to manually request upstream DDoS mitigation service initiation from a DOTS server. When an organization has reason to believe that it is under active attack, authorized personnel may utilize the Web portal to manually initiate a DOTS client mitigation request to one or more DOTS servers. DDoS mitigation service may be terminated manually via a DOTS mitigation service termination request through the Web portal when it has been determined that the DDoS attack has ended.

In this use-case, the organization targeted for attack does not possess any automated or operator-assisted mechanisms for DDoS attack

detection, classification, traceback, or mitigation; the existence of an attack has been inferred manually, and the organization is requesting upstream mitigation assistance. This can theoretically be an inter- or intra-domain use-case, but is more typically an inter-domain scenario.

(a) A DDoS attack is initiated against online properties of an organization have access to a Web portal which incorporates DOTS client functionality and can generate DOTS mitigation service requests upon demand.

(b) Authorized personnel utilize the Web portal to send a DOTS mitigation service initiation request to one or more upstream transit networks, peer networks, or overlay MSSP networks, either directly or via intermediate DOTS relays residing upon the requesting organization's network, the upstream mitigation provider's network, or both. This DOTS mitigation service initiation request is manually triggered by personnel of the requesting organization when it is judged that the organization is under DDoS attack (the mechanism by which this process takes place is beyond the scope of this document).

(c) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been provisioned to honor requests from the Web portal, and initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

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

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the Web portal.

(f) While DDoS mitigation services are active, the Web portal may optionally regularly transmit manually-triggered DOTS mitigation efficacy updates to the relevant DOTS servers.

(g) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to their respective DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(h) The DOTS servers transmit a DOTS mitigation status update to the Web portal indicating that the DDoS attack has ceased.

(i) The Web portal transmits a manually-triggered DOTS mitigation service termination request to the DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(j) The Web portal transmits a manually-triggered DOTS mitigation service termination request to the DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(k) The DOTS servers transmit a DOTS mitigation status update to the Web portal indicating that DDoS mitigation services have been terminated.

(l) The Web portal transmits a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.6. Manual Mobile Device Application Request to Upstream Mitigator

An application for mobile devices such as smartphones and tablets which incorporates DOTS client capabilities has been made available to authorized personnel of an organization. When the organization has reason to believe that it is under active DDoS attack, authorized personnel may utilize the mobile device application to manually initiate a DOTS client mitigation request to one or more DOTS servers in order to initiate upstream DDoS mitigation services. DDoS mitigation service may be terminated manually via a DOTS mitigation service termination request initiated through the mobile device application when it has been determined that the DDoS attack has ended.

This use-case is similar to the one described in Section 3.1.5; the difference is that a mobile application provided by the DDoS mitigation service provider is used to request upstream attack mitigation assistance. This can theoretically be an inter- or intra-domain use-case, but is more typically an inter-domain scenario.

(a) A DDoS attack is initiated against online properties of an organization have access to a Web portal which incorporates DOTS client functionality and can generate DOTS mitigation service requests upon demand.

(b) Authorized personnel utilize the mobile application to send a DOTS mitigation service initiation request to one or more DOTS servers residing on the same network as the targeted Internet properties, one or more upstream transit networks, peer networks, or overlay MSSP networks, either directly or via intermediate DOTS relays residing upon the requesting

organization's network, the upstream mitigation provider's network, or both. This DOTS mitigation service initiation request is manually triggered by personnel of the requesting organization when it is judged that the organization is under DDoS attack (the mechanism by which this process takes place is beyond the scope of this document).

(c) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been provisioned to honor requests from the mobile application, and initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

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

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the mobile application.

(f) While DDoS mitigation services are active, the mobile application may optionally regularly transmit manually-triggered DOTS mitigation efficacy updates to the relevant DOTS servers.

(g) When the upstream DDoS mitigators determine that the DDoS attack has ceased, they indicate this change in status to their respective DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(h) The DOTS servers transmit a DOTS mitigation status update to the mobile application indicating that the DDoS attack has ceased.

(i) The mobile application transmits a manually-triggered DOTS mitigation service termination request to the DOTS servers (the mechanism by which this process takes place is beyond the scope of this document).

(j) The DOTS servers terminate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(k) The DOTS servers transmit a DOTS mitigation status update to the mobile application indicating that DDoS mitigation services have been terminated.

(l) The mobile application transmits a DOTS mitigation termination status acknowledgement to the DOTS servers.

3.1.7. Unsuccessful Automatic or Operator-Assisted CPE or PE Mitigators Request Upstream DDoS Mitigation Services

One or more CPE or PE mitigators with DOTS client capabilities may be configured to signal to one or more DOTS servers in order to request upstream DDoS mitigation service initiation during an attack when DDoS attack volumes and/or attack characteristics exceed the capabilities of such CPE mitigators. DDoS mitigation service may be terminated either automatically or manually via a DOTS mitigation service termination request initiated by the mitigator when it has been determined that the DDoS attack has ended.

This can theoretically be an inter- or intra-domain use-case, but is more typically an inter-domain scenario.

(a) A DDoS attack is initiated against online properties of an organization which has deployed DOTS-client-capable DDoS mitigators.

(b) CPE or PE DDoS mitigators detect, classify, and begin mitigating the DDoS attack.

(c) CPE or PE DDoS mitigators determine that their capacity and/or capability to mitigate the DDoS attack is insufficient, and utilize their DOTS client functionality to send a DOTS mitigation service initiation request to one or more DOTS servers residing on one or more upstream transit networks, peer networks, or overlay MSSP networks. This DOTS mitigation service initiation request may be automatically initiated by the CPE or PE DDoS mitigators, or may be manually triggered by personnel of the requesting organization in response to an alert from the mitigators (the mechanism by which this process takes place is beyond the scope of this document).

(d) The DOTS servers which receive the DOTS mitigation service initiation requests determine that they have been configured to honor requests from the requesting CPE or PE mitigators, and attempt to initiate situationally-appropriate DDoS mitigation service on their respective networks (the mechanism by which this process takes place is beyond the scope of this document).

(e) The DDoS mitigators on the upstream network report back to the DOTS servers 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 (the mechanism by which this process takes place is beyond the scope of this document).

(f) The DOTS servers transmit a DOTS service status message to the requesting CPE or PE mitigators indicating that upstream DDoS mitigation service cannot be initiated as requested.

(g) The CPE or PE mitigators may optionally regularly re-transmit DOTS mitigation status request messages to the relevant DOTS servers until acknowledgement that mitigation services have been initiated.

(h) The CPE or PE mitigators may optionally transmit a DOTS mitigation service initiation request to DOTS servers 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 CPE or PE mitigators determine that the DDoS attack volume has decreased to a level and/or complexity which they themselves can successfully mitigate; the DDoS attack has ceased; or manual intervention by personnel of the requesting organization has taken place.

3.2. Ancillary Use Cases

3.2.1. Auto-registration of DOTS clients with DOTS servers

An additional benefit of DOTS is that by utilizing agreed-upon authentication mechanisms, DOTS clients can automatically register for DDoS mitigation service with one or more upstream DOTS servers. The details of such registration are beyond the scope of this document.

3.2.2. Auto-provisioning of DDoS countermeasures

The largely manual tasks associated with provisioning effective, situationally-appropriate DDoS countermeasures is a significant barrier to providing/obtaining DDoS mitigation services for both mitigation providers and mitigation recipients. Due to the 'self-descriptive' nature of DOTS registration messages and mitigation requests, the implementation and deployment of DOTS has the potential to automate countermeasure selection and configuration for DDoS mitigators. The details of such provisioning are beyond the scope of this document.

This can theoretically be an inter- or intra-domain use-case, but is more typically an inter-domain scenario.

3.2.3. Informational DDoS attack notification to interested and authorized third parties

In addition to its primary role of providing a standardized, programmatic approach to the automated and/or operator-assisted request of DDoS mitigation services and providing status updates of those mitigations to requesters, DOTS may be utilized to notify security researchers, law enforcement agencies, regulatory bodies, etc. of DDoS attacks against attack targets, assuming that organizations making use of DOTS choose to share such third-party notifications, in keeping with all applicable laws, regulations, privacy and confidentiality considerations, and contractual agreements between DOTS users and said third parties.

This is an inter-domain scenario.

4. Security Considerations

DOTS is at risk from three primary attacks: DOTS agent impersonation, traffic injection, and signaling blocking. The DOTS protocol MUST be designed for minimal data transfer to address the blocking risk.

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

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

5. IANA Considerations

No IANA considerations exist for this document at this time.

6. Acknowledgments

TBD

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

7.2. Informative References

[APACHE] "Apache mod_security", <https://www.modsecurity.org>.

[I-D.ietf-dots-requirements] Mortensen, A., Moskowitz, R., and T. Reddy, "DDoS Open Threat Signaling Requirements", draft-ietf-dots-requirements-00 (work in progress), October 2015.

[RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP 165, RFC 6335, DOI 10.17487/RFC6335, August 2011, <http://www.rfc-editor.org/info/rfc6335>.

[RRL] "BIND RRL", <https://deepthought.isc.org/article/AA-00994/0/Using-the-Response-Rate-Limiting-Feature-in-BIND-9.10.html>.

Authors' Addresses

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

Email: rdobbins@arbor.net

Stefan Fouant
Corero Network Security

Email: Stefan.Fouant@corero.com

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

Phone: +1 514-452-2160
Email: daniel.migault@ericsson.com

Robert Moskowitz
HTT Consulting
Oak Park, MI 48237
USA

Email: rgm@labs.htt-consult.com

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

Phone: +44 791 763 5384
Email: nteague@verisign.com

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

Email: Frank.xialiang@huawei.com