DOTS
Internet Draft
Intended status: Standard Track
Expires: April 2016

T. Fu
Huawei
D. Zhang
Alibaba
L. Xia
M. Li
Huawei

October 19, 2015

IPFIX IE Extensions for DDoS Attack Detection
draft-fu-dots-ipfix-extension-00.txt

Status of this Memo

This Internet-Draft will expire on April 19, 200916.

Abstract

Although most of the existing IP Flow Information Export (IPFIX) Information Elements (IEs) are useful for network security inspection, there are still some gaps existing to identify a number of categories of the attacks. To fill in the gaps, this document defines some new IPFIX IEs and describes their formats for inspecting network security.

Table of Contents

1. Introduction

As network security issues arising dramatically nowadays, network administrators are eager to detect and identify attacks as early as possible, generate countermeasures with high agility. Due to the enormous amount of network attack types, metrics useful for attack detection are also enormous. Moreover, attacking methods are evolved rapidly, which brings challenges to designing detection mechanism.

The IPFIX Protocol [RFC7011] defines a generic exchange mechanism for flow information and events. It supports source-triggered exporting of information via the push model approach. The IPFIX Information Model [IPFIX-IANA] defines a list of standard Information Elements (IEs) which can be carried by the IPFIX protocol. The IPFIX requirement [RFC3917] points out that one of the target applications of IPFIX is attack and intrusion detection. Although the existing standard IEs provide a rich source of data for security inspection by checking the status/events of the traffic, there are still some gaps existing to identify a number of categories of the attacks. The scanty information will result in an inaccurate analysis and slowing down the effective response towards network attacks. More detailed gap analysis is given in the following section.

This document presents the IPFIX IEs which are available for the network attacks detection, some of them are the new defined IPFIX IEs and their formats are specified. The wise utilization of these IEs will improve the network security and will support the offline analysis of data from different operators in the future with minimal resource consumption.

This document is structured as following: Section 3 discusses the connection sampling mechanism and introduces the new IPFIX IEs derived from relevant use cases. Section 4 describes how to use these IEs to detect specific DDoS attacks.

2. Conventions used in this document

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].

2.1. Terminology

IPFIX-specific terminology (Information Element, Template, Template Record, Options Template Record, Template Set, Collector, Exporter,

Data Record, etc.) used in this document is defined in Section 2 of [RFC7011]. As in [RFC7011], these IPFIX-specific terms have the first letter of a word capitalized.

3. Connection Sampling and new IEs

3.1. Packet Sampling vs Connection Sampling

Packet sampling is a widely used method to select packets from network traffic for reporting. Its selection operations include random selection, deterministic selection, hash-based selection and so on. Although it is easy and efficient, it still has a number of limitations:

- Several research projects [N. DUFFIELD, 2003], [D. BRAUCKHOFF 2006] show that packet sampling impacts larger on small flows (with only few packets) due to the smaller sampling probability compared to larger flows, unfortunately attacks such as SYN-Flood, ACK-Flood all have small flow characteristics, which means that packet sampling may impair the detection performance for small flow based DDoS attacks;

- Although the communication is 2-way between source and destination, current packet sampling is applied independently in each direction, which leads to difficulties correlating the statistic of both sides, despite that those metrics are essential indicators in detecting attacks such as SNMP/DNS Reflected Amplification (i.e. where there are not much or even no traffic in the opposite direction of the attacking flow);

- Today’s packet sampling cannot provide detailed information of traffic between communication peers, which makes it impossible to distinguish some of the attacks, such as IP fragment attack and Slowloris HTTP attack, from ordinary traffic.

As a consequence from the above analysis, a layer 4 connection oriented sampling method is more suitable for the security application: Rather than sampling a small part of packets in the traffic between the communication peers, the connection sampling records all TCP/UDP connection packets (including packets during connection setup and close phase if there is) between them once that connection is selected to be sampled. Furthermore, several new IPFIX IEs are proposed in this document to represent the telemetry information that can obtain via this method.

3.2. Use Cases for New IEs

In this section, several use cases are discussed to identify the requirements where new IEs are desirable for the network attacks detection.

3.2.1. Upstream/Downstream Counters

Take ICMP attack as an example, ICMP flow model has features such as the ICMP Echo/Echo Reply dominate the whole traffic flow, ICMP packet interval is usually not too short (normally 1 pkt/s). Usually, the normal ratio between ICMP echo to ICMP echo reply packets is around 1:1. When a DDOS attack happens, a sudden burst of messages from different sources to a destination endpoint can be detected. In turn, the ratio between echo and echo reply packets will be significantly biased from the normal ratio, i.e., exceed 20:1. So, the proper way to distinguish an attack from the normal communication is to check this ratio.

However, the current IPFIX IEs for ICMP contain the ICMP type and code for both IPv4 and IPv6 only for a single ICMP packet rather than statistical property of the ICMP session. Further metrics like the cumulated sum of various counters should be calculated based on sampling method defined by the Packet SAMPling (PSAMP) protocol [RFC5477]. Similar problems occur in TCP, UDP, SNMP and DNS attacks. It would be useful to calculate the number of the upstream and downstream packets for one connection separately over time in order to detect the anomalies of the network. For ICMP attack, a more generic approach is to define two basic metrics icmpEchoCount and icmpEchoReplyCount as new IPFIX IEs to represent the cumulated upstream and downstream packets counter within a ICMP connection. Similar new IE definitions of pktUpstreamCount, pktDownstreamCount, octetUpstreamCount, and octetDownstreamCount are applied to the TCP, UDP, SNMP and DNS attacks.

3.2.2. Fragment statistic

Fragment attack employs unexpected formats of fragmentation, e.g. without last fragment or incorrect fragment offset[RFC791], which result in errors such as fragmentation buffer overrun and fragment overlapped. Existing IPFIX fragmentation metrics includes fragmentOffset, fragmentIdentification, fragmentFlags, which only indicate the attributes of a single fragment, and are not suitable for attack detection. Instead, the network attack should be observed based upon a historic, integrated view of fragmented packets of a connection. For instances, if more than 500 out of 1000 fragmented

packets have fragment errors, it is likely that a fragment attack happens.

Therefore, a number of new IEs associated with fragment statistics are proposed as follows:

- fragmentIncompleteCount: The completeness of fragmented packets of the same connection should be checked, and this metric is proposed to count the incomplete events;

- fragmentFirstTooShortCount: Attacker might intent to exclude destination port from the first fragment so as to bypass detection from firewall. This metric is proposed to indicate the number of the invalid first fragments in the observed connection;

- fragmentOffestErrorCount: This metric is proposed to count the number of fragments with offset error, and the value can be used to indicate attack occurs;

- fragmentFlagErrorCount: This metric is proposed to detect early whether the fragment flags are incorrectly set on purpose.

3.2.3. Extended Value of FlowEndReason

Refer to [IPFIX-IANA], there are 5 defined reasons for Flow termination, with values ranging from 0x01 to 0x05:

0x01: idle timeout

0x02: active timeout

0x03: end of Flow detected

0x04: forced end

0x05: lack of resources

There is an additional reason caused by state machine anomaly. When FIN/SYN is sent, but no ACK is replied after a waiting timeout, the existing five reasons do not match this case. Therefore, a new value is proposed to extend the FlowEndReason, which is 0x06: protocol exception timeout.

3.3. Definition of New IEs

The following is the table of all the IEs that a device would need to export for attack statistic analysis. The formats of the IEs and

the IPFIX IDs are listed below, as well as their descriptions. Some of the IEs are already defined in [IPFIX-IANA]. While a number of new IE's IDs are not assigned yet, their explanations are presented in the previous sections. The recommended registrations to IANA are described in the IANA considerations section.

Field Name	Size (bits)	IANA IPFIX ID	Description
sourceIPv4Address	32	8	Source IPv4 Address
destinationIPv4Address	32	12	Destination IPv4 Address
sourceTransportPort	16	7	Source Port
destinationTransportPort	16	11	Destination port
protocolIdentifier	8	4	Transport protocol
packetDeltaCount	64	2	The number of incoming packets since the previous report (if any) for this Flow at the Observation Point
pktUpstreamCount	64	TBD	Upstream packet counter
pktDownstreamCount	64	TBD	Downstream packet counter
octetUpstreamCount	64	TBD	Upstream octet counter
octetDownstreamCount	64	TBD	Downstream octet counter
tcpSynTotalCount	64	218	The total number of packets of this Flow with TCP "Synchronize sequence numbers" (SYN) flag set
tcpFinTotalCount	64	219	The total number of packets of this Flow with TCP "No more data from sender" (FIN)

			flag set
tcpRstTotalCount	64	220	The total number of packets of this Flow with TCP "Reset the connection" (RST) flag set.
tcpPshTotalCount	64	221	The total number of packets of this Flow with TCP "Push Function" (PSH) flag set.
tcpAckTotalCount	64	222	The total number of packets of this Flow with TCP "Acknowled gment field significant" (ACK) flag set.
tcpUrgTotalCount	64	223	The total number of packets of this Flow with TCP "Urgent Pointer field significant" (URG) flag set.
tcpControlBits	8	6	TCP control bits observed for packets of this Flow
flowEndReason	8	136	The reason for Flow termination
minimumIpTotalLength	64	25	Length of the smallest packet observed for this Flow

maximumIpTotalLength	64	26	Length of the largest packet observed for this Flow
flowStartSeconds	dateTimeSeconds	150	The absolute timestamp of the first packet of this Flow
flowEndSeconds	dateTimeSeconds	151	The absolute timestamp of the last packet of this Flow
flowStartMilliseconds	dateTimeMilliseconds	152	The absolute timestamp of the first packet of this Flow
flowEndMilliseconds	dateTimeMilliseconds	153	The absolute timestamp of the last packet of this Flow
flowStartMicroseconds	dateTimeMicroseconds	154	The absolute timestamp of the first packet of this Flow
flowEndMicroseconds	dateTimeMicroseconds	155	The absolute timestamp of the last packet of this Flow
fragmentFlags	8	197	Fragmentation properties indicated by flags in the IPv4 packet header or the IPv6 Fragment header, respectively
fragmentPacketDeltaCount	32	TBD	Counter of session fragments
fragmentFirstTooShort	32	TBD	Number of

DeltaCount			packets with first fragment too short
fragmentFlagErrorDeltaCount	32	TBD	Number of fragments with flag error
icmpTypeIPv4	8	176	Type of the IPv4 ICMP message
icmpCodeIPv4	8	177	Code of the IPv4 ICMP message
icmpTypeIPv6	8	178	Type of the IPv6 ICMP message
icmpCodeIPv6	8	179	Code of the IPv6 ICMP message
icmpEchoDeltaCount	32	TBD	The number fo ICMP echo.
icmpEchoReplyDeltaCount	32	TBD	The number of ICMP echo reply.
selectorAlgorithm	16	304	This Information Element identifies the packet selection methods (e.g., Filtering, Sampling) that are applied by the Selection Process.
samplingPacketInterval	32	305	The number of packets that are consecutively sampled
samplingPacketSpace	32	306	The number of packets between two "samplingPacketInterval"s.
octetVariance	64	TBD	IP packet byte variance

			statistic
tcpControlStateBits	16	TBD	tcp states
flowSessionEndMilliseconds	64	TBD	the absolute timestamp of the first FIN or RST packet of this Flow
tcpPayloadOctetTotalCount	64	TBD	tcp payload statistics, it is equal to the ACK's window value subtract INIT's window value
tcpOutoforderTotalCount	64	TBD	out of order packets statistic
flowTimeIntervalVariance	64	TBD	the interval time variance between upstream and downstream traffic of a flow
flowTimeInterval	64	TBD	the interval time between upstream and downstream traffic of a flow
serverResponseTime	64	TBD	the response time of a server
clientResponseTime	64	TBD	the response time of a client
sessionResponseTime	64	TBD	the response time of a session

Table 1: Information Element Table

4. Application of the New IEs for Attack Detection

This section presents a number of examples to help for the easy understanding of the application of these new IEs for attack detection.

4.1. Use of Upstream/Downstream Counters to Detect ICMP Attack

According to previous analysis, the template for detecting ICMP attack should at least contain IEs shown in Table 2.

Set ID = 2	Length = 24 octets
Template ID TBD	Field Count = 10
0 sourceIPv4Address	Field Length = 4
0 destinationIPv4Address	Field Length = 4
0 protocolIdentifier	Field Length = 1
0 packetDeltaCount	Field Length = 8
0 protocolIdentifier	Field Length = 1
0 packetDeltaCount	Field Length = 8
0 pktUpstreamCount	Field Length = 4
0 pktDownstreamCount	Field Length = 4
0 flowStartSeconds	Field Length = 4
0 flowEndSeconds	Field Length = 4

Table 2: Template example for detecting ICMP attack

An example of the actual ICMP event data record is shown below in a readable form as below:

{sourceIPv4Address = 192.168.0.101, destinationIPv4Address = 192.168.0.201, protocolIdentifier = 1, packetDeltaCount = 3000, icmpEchoCount = 2880, icmpEchoRelayCount = 120, flowStartSeconds = 100, flowEndSeconds = 200}

protocolIdentifier = 1 represents the ICMP proptocol. There are 30 ICMP messages transmited per second. Tthe ICMP Echo to ICMP Echo Reply packet ratio is 24:1, which indicates a high possibility of ICMP attack.

4.2. Fragment Attack

The template for detecting fragment attack should at least contain IEs shown in Table 3. It requires the observation point to trace complete fragmented packet and accumulate the errors.

Set ID = 2	Length = 24 octets
Template ID TBD	Field Count = 10
0 sourceIPv4Address	Field Length = 4
0 destinationIPv4Address	Field Length = 4
0 protocolIdentifier	Field Length = 1
0 fragmentIncompleteCount	Field Length = 4
0 fragmentFirstTooShortCount	Field Length = 4
0 fragmentOffestErrorCount	Field Length = 4
0 fragmentFlagErrorCount	Field Length = 4
0 flowStartSeconds	Field Length = 4
0 flowEndSeconds	Field Length = 4

Table 3: Template example for detecting fragment attack

An example of the actual fragment attack record is shown below in a readable form as below:

{sourceIPv4Address = 192.168.0.101, destinationIPv4Address = 192.168.0.201, protocolIdentifier = 1,fragmentIncompleteCount = 0, fragmentFirstTooShortCount = 0, fragmentOffestErrorCount = 3000, fragmentFlagErrorCount = 0, flowStartSeconds = 100, flowEndSeconds = 200}

In this case, fragment offset errors are used to exhaust resource at the receiver.

5. Security Considerations

No additional security considerations are introduced in this document. The same security considerations as for the IPFIX protocol [RFC7011] apply.

6. IANA Considerations

The following information elements are requested from IANA IPFIX registry.

Name : pktUpstreamCount

Description: The number of the upstream packets for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: TBD

Name: pktDownstreamCount

Description: The number of the downstream packets for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: TBD

Name: octetUpstreamCount

Description: The total number of octets in upstream packets for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point. The number of octets includes IP header(s) and IP payload.

Abstract Data Type: unsigned64

Data Type Semantics: TBD

Name : octetDownstreamCount

Description: The total number of octets in downstream packets for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point. The number of octets includes IP header(s) and IP payload.

Abstract Data Type: unsigned64

Data Type Semantics: TBD

Name: fragmentPacketDeltaCount

Description: This Information Element is the counter of session fragments.

Abstract Data Type: unsigned32

Data Type Semantics: TBD

Name: fragmentFirstTooShortDeltaCount

Description: This Information Element indicates the number of packets with first fragment too short.

Abstract Data Type: unsigned32

Data Type Semantics: TBD

Name: fragmentFlagErrorDeltaCount

Description: This Information Element specifies number of fragments with offset error.When the DF bit and MF bit of the fragment flag are set in the same fragment, there is an error at the fragment flag.

Abstract Data Type: unsigned32

Data Type Semantics: TBD

Name: octetVariance

Description: IP packet byte variance statistic.

Abstract Data Type: unsigned64

Data Type Semantics: quantity

Name: tcpControlStateBits

Description: tcp states.

Abstract Data Type: unsigned16

Data Type Semantics: flags

Name: icmpEchoDeltaCount

Description: icmp Echo packets.

Abstract Data Type: unsigned32

Data Type Semantics: deltaCounter

Name: icmpEchoReplyDeltaCount

Description: icmp Echo Reply packets.

Abstract Data Type: unsigned32

Data Type Semantics: deltaCounter

Name: flowSessionEndMilliseconds

Description: the absolute timestamp of the first FIN or RST packet of this flow.

Abstract Data Type: dateTimeMilliseconds

Data Type Semantics: default

Name: tcpPayloadOctetTotalCount

Description: tcp payload statistics, it is equal to the ACK's window value subtract INIT's window value.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Name: tcpOutoforderTotalCount

Description: out of order packets statistic.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Name: flowTimeIntervalVariance

Description: the interval time variance between upstream and downstream.

Abstract Data Type: unsigned64

Data Type Semantics: quantity

Name: flowTimeInterval

Description: the interval time between upstream and downstream.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Name: flowTimeInterval

Description: the interval time between upstream and downstream.

Abstract Data Type: unsigned64

Data Type Semantics: quantity

Name: serverResponseTime

Description: the response time of a server.

Abstract Data Type: unsigned16

Data Type Semantics: quantity

Name: clientResponseTime

Description: the response time of a client.

Abstract Data Type: unsigned16

Data Type Semantics: quantity

Name: sessionResponseTime

Description: the response time of a session.

Abstract Data Type: unsigned16

Data Type Semantics: quantity

A new values is added to FlowEndReason:

0x06: protocol exception timeout

The flow was terminated due to protocol state machine anomaly and unexpected timeout.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

[RFC3917] Quittek, J., Zseby, T., Claise, B., Zander, S., "Requirements for IP Flow Information Export (IPFIX)", RFC 3917, October 2004.

7.2. Informative References

[IPFIX-IANA]

IANA, "IPFIX Information Elements registry",

<http://www.iana.org/assignments/ipfix>.

[D. BRAUCKHOFF 2006]

Daniela Brauckhoff, Bernhard Tellenbach, Arno Wagner, Martin May, and Anukool Lakhina. 2006. Impact of packet sampling on anomaly detection metrics. In Proceedings of the 6th ACM SIGCOMM conference on Internet measurement (IMC '06). ACM, New York, NY, USA, 159-164.

[N. DUFFIELD, 2003]

DUFFIELD, N., LUND, C., AND THORUP, M., Estimating Flow Distributions from Sampled Flow Statistics. In ACM SIGCOMM (Karlsruhe, August 2003).

8. Acknowledgments

The authors would thank Danping He and Yibo Zhang for their great help during the initial period of this draft.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Tianfu Fu
Huawei
Q11, Huanbao Yuan, 156 Beiqing Road, Haidian District
Beijing 100095
China

Email: futianfu@huawei.com

DaCheng Zhang
Alibaba

Email: Dacheng.zdc@alibaba-inc.com

Liang Xia (Frank)
Huawei

101 Software Avenue, Yuhuatai District
Nanjing, Jiangsu 210012
China

Email: Frank.xialiang@huawei.com

Min Li
Huawei

Huawei Technologies Duesseldorf GmbH, European Research Center,
Riesstr. 25, 80992 Muchen, Germany
Email: l.min@huawei.com

DOTS
Internet-Draft
Intended status: Informational
Expires: April 21, 2016

A. Mortensen
Arbor Networks, Inc.
R. Moskowitz
HTT Consulting
T. Reddy
Cisco Systems, Inc.

October 19, 2015

DDoS Open Threat Signaling Requirements
draft-ietf-dots-requirements-00

Abstract

This document defines the requirements for the DDoS Open Threat Signaling (DOTS) protocols coordinating attack response against Distributed Denial of Service (DDoS) attacks.

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 April 21, 2016.

Copyright Notice

Table of Contents

1. Introduction

1.1. Overview

Distributed Denial of Service (DDoS) attacks continue to plague networks around the globe, from Tier-1 service providers on down to enterprises and small businesses. Attack scale and frequency similarly have continued to increase, thanks to software vulnerabilities leading to reflection and amplification attacks. Once staggering attack traffic volume is now the norm, and the impact of larger-scale attacks attract the attention of international press agencies.

The higher profile and greater impact of contemporary DDoS attacks has led to increased focus on coordinated attack response. Many institutions and enterprises lack the resources or expertise to operate on-premise attack prevention solutions themselves, or simply find themselves constrained by local bandwidth limitations. To address such gaps, security service providers have begun to offer on-demand traffic scrubbing services. Each service offers its own interface for subscribers to request attack mitigation, tying subscribers to proprietary implementations while also limiting the subset of network elements capable of participating in the attack response. As a result of incompatibility across services, attack

response may be fragmentary or otherwise incomplete, leaving key players in the attack path unable to assist in the defense.

There are many ways to respond to an ongoing DDoS attack, some of them better than others, but the lack of a common method to coordinate a real-time response across layers and network domains inhibits the speed and effectiveness of DDoS attack mitigation.

DOTS was formed to address this lack. The DOTS protocols are therefore not concerned with the form of response, but rather with communicating the need for a response, supplementing the call for help with pertinent details about the detected attack. To achieve this aim, the protocol must permit the DOTS client to request or withdraw a request for coordinated mitigation; to set the scope of mitigation, restricted to the client's network space; and to supply summarized attack information and additional hints the DOTS server elements can use to increase the accuracy and speed of the attack response.

The protocol must also continue to operate even in extreme network conditions. It must be resilient enough to ensure a high probability of signal delivery in spite of high packet loss rates. As such, elements should be in regular, bidirectional contact to measure peer health, provide mitigation-related feedback, and allow for active mitigation adjustments.

Lastly, the protocol must take care to ensure the confidentiality, integrity and authenticity of messages passed between peers to prevent the protocol from being repurposed to contribute to the very attacks it's meant to deflect.

Drawing on the DOTS use cases [I-D.ietf-dots-use-cases] for reference, this document details the requirements for protocols achieving the DOTS goal of standards-based open threat signaling.

1.2. Terminology

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

The following terms are used to define relationships between elements, the data they exchange, and methods of communication among them:

attack telemetry: collected network traffic characteristics defining the nature of a DDoS attack.

mitigation: A defensive response against a detected DDoS attack, performed by an entity in the network path between attack sources and the attack target, either through inline deployment or some form of traffic diversion. The form mitigation takes is out of scope for this document.

mitigator: A network element capable of performing mitigation of a detected DDoS attack.

DOTS client: A DOTS-aware network element requesting attack response coordination with another DOTS-aware element, with the expectation that the remote element is capable of helping fend off the attack against the client.

DOTS server: A DOTS-aware network element handling and responding to messages from a DOTS client. The DOTS server MAY enable mitigation on behalf of the DOTS client, if requested, by communicating the DOTS client's request to the mitigator and relaying any mitigator feedback to the client. A DOTS server may also be a mitigator.

DOTS relay: A DOTS-aware network element positioned between a DOTS server and a DOTS client. A DOTS relay receives messages from a DOTS client and relays them to a DOTS server, and similarly passes messages from the DOTS server to the DOTS client.

DOTS agents: A collective term for DOTS clients, servers and relays.

signal channel: A bidirectional, mutually authenticated communication layer between DOTS agents characterized by resilience even in conditions leading to severe packet loss, such as a volumetric DDoS attack causing network congestion.

DOTS signal: A concise authenticated status/control message transmitted between DOTS agents, used to indicate client's need for mitigation, as well as to convey the status of any requested mitigation.

heartbeat: A keep-alive message transmitted between DOTS agents over the signal channel, used to measure peer health. Heartbeat functionality is not required to be distinct from signal.

client signal: A message sent from a DOTS client to a DOTS server over the signal channel, possibly traversing a DOTS relay, indicating the DOTS client's need for mitigation, as well as the scope of any requested mitigation, optionally including detected attack telemetry to supplement server-initiated mitigation.

server signal: A message sent from a DOTS server to a DOTS client over the signal channel. Note that a server signal is not a response to client signal, but a DOTS server-initiated status message sent to the DOTS client, containing information about the status of any requested mitigation and its efficacy.

data channel: A secure communication layer between client and server used for infrequent bulk exchange of data not easily or appropriately communicated through the signal channel under attack conditions.

blacklist: a list of source addresses or prefixes from which traffic should be blocked.

whitelist: a list of source addresses or prefixes from which traffic should always be allowed, regardless of contradictory data gleaned in a detected attack.

2. Requirements

This section describes the required features and characteristics of the DOTS protocols. The requirements are informed by the use cases described in [I-D.ietf-dots-use-cases].

DOTS must at a minimum make it possible for a DOTS client to request a DOTS server’s aid in mounting a coordinated defense against a detected attack, by signaling inter- or intra-domain using the DOTS protocol. DOTS clients should similarly be able to withdraw aid requests arbitrarily. Regular feedback between DOTS client and server supplement the defensive alliance by maintaining a common understanding of DOTS peer health and activity. Bidirectional communication between DOTS client and server is therefore critical.

Yet the DOTS protocol must also work with a set of competing operational goals. On the one hand, the protocol must be resilient under extremely hostile network conditions, providing continued contact between DOTS agents even as attack traffic saturates the link. Such resiliency may be developed several ways, but characteristics such as small message size, asynchronous, redundant message delivery and minimal connection overhead (when possible given local network policy) with a given network will tend to contribute to the robustness demanded by a viable DOTS protocol.

On the other hand, DOTS must have adequate message confidentiality, integrity and authenticity to keep the protocol from becoming another vector for the very attacks it’s meant to help fight off. The DOTS client must be authenticated to the DOTS server, and vice versa, for DOTS to operate safely, meaning the DOTS agents must have a way to

negotiate and agree upon the terms of protocol security. Attacks against the transport protocol should not offer a means of attack against the message confidentiality, integrity and authenticity.

The DOTS server and client must also have some common method of defining the scope of any mitigation performed by the mitigator, as well as making adjustments to other commonly configurable features, such as listen ports, exchanging black- and white-lists, and so on.

Finally, DOTS should provide sufficient extensibility to meet local, vendor or future needs in coordinated attack defense, although this consideration is necessarily superseded by the other operational requirements.

2.1. General Requirements

G-001 Interoperability: DOTS's objective is to develop a standard mechanism for signaling detected ongoing DDoS attacks. That objective is unattainable without well-defined specifications for any protocols or data models emerging from DOTS. All protocols, data models and interfaces MUST be detailed enough to ensure interoperable implementations.

G-002 Extensibility: Any protocols or data models developed as part of DOTS MUST be designed to support future extensions. Provided they do not undermine the interoperability and backward compatibility requirements, extensions are a critical part of keeping DOTS adaptable to changing operational and proprietary needs to keep pace with evolving DDoS attack methods.

G-003 Resilience: The signaling protocol MUST be designed to maximize the probability of signal delivery even under the severely constrained network conditions imposed by the attack traffic. The protocol SHOULD be resilient, that is, continue operating despite message loss and out-of-order or redundant signal delivery.

G-004 Bidirectionality: To support peer health detection, to maintain an open signal channel, and to increase the probability of signal delivery during attack, the signal channel MUST be bidirectional, with client and server transmitting signals to each other at regular intervals, regardless of any client request for mitigation.

G-005 Sub-MTU Message Size: To avoid message fragmentation and the consequently decreased probability of message delivery, signaling protocol message size MUST be kept under signaling path Maximum Transmission Unit (MTU), including the byte overhead of any

encapsulation, transport headers, and transport- or message-level security.

G-006 Message Integrity: DOTS protocols MUST take steps to protect the confidentiality, integrity and authenticity of messages sent between client and server. While specific transport- and message-level security options are not specified, the protocols MUST follow current industry best practices for encryption and message authentication.

In order for DOTS protocols to remain secure despite advancements in cryptanalysis, DOTS agents MUST be able to negotiate the terms and mechanisms of protocol security, subject to the interoperability and signal message size requirements above.

G-007 Message Replay Protection: In order to prevent a passive attacker from capturing and replaying old messages, DOTS protocols MUST provide a method for replay detection, such as including a timestamp or sequence number in every heartbeat and signal sent between DOTS agents.

G-008 Bulk Data Exchange: Infrequent bulk data exchange between DOTS client and server can also significantly augment attack response coordination, permitting such tasks as population of black- or white-listed source addresses; address group aliasing; exchange of incident reports; and other hinting or configuration supplementing attack response.

As the resilience requirements for DOTS mandate small signal message size, a separate, secure data channel utilizing an established reliable protocol SHOULD be used for bulk data exchange. The mechanism for bulk data exchange is not yet specified, but the nature of the data involved suggests use of a reliable, adaptable protocol with established and configurable conventions for authentication and authorization.

2.2. Operational requirements

OP-001 Use of Common Transports: DOTS MUST operate over common standardized transport protocols. While the protocol resilience requirement strongly RECOMMENDS the use of connectionless protocols, in particular the User Datagram Protocol (UDP) [RFC0768], use of a standardized, connection-oriented protocol like the Transmission Control Protocol (TCP) [RFC0793] MAY be necessary due to network policy or middleware limitations.

OP-002 Peer Mutual Authentication: The client and server MUST authenticate each other before a DOTS session is considered

active. The method of authentication is not specified, but should follow current industry best practices with respect to any cryptographic mechanisms to authenticate the remote peer.

OP-003 Session Health Monitoring: The client and server MUST regularly send heartbeats to each other after mutual authentication in order to keep the DOTS session open. A session MUST be considered active until a client or server explicitly ends the session, or either DOTS agent fails to receive heartbeats from the other after a mutually negotiated timeout period has elapsed.

OP-004 Mitigation Capability Opacity: DOTS is a threat signaling protocol. The server and mitigator MUST NOT make any assumption about the attack detection, classification, or mitigation capabilities of the client. While the server and mitigator MAY take hints from any attack telemetry included in client signals, the server and mitigator cannot depend on the client for authoritative attack classification. Similarly, the mitigator cannot assume the client can or will mitigate attack traffic on its own.

The client likewise MUST NOT make any assumptions about the capabilities of the server or mitigator with respect to detection, classification, and mitigation of DDoS attacks. The form of any attack response undertaken by the mitigator is not in scope.

OP-005 Mitigation Status: DOTS clients MUST be able to request or withdraw a request for mitigation from the DOTS server. The DOTS server MUST acknowledge a DOTS client’s request to withdraw from coordinated attack response in subsequent signals, and MUST cease mitigation activity as quickly as possible. However, a DOTS client rapidly toggling active mitigation may result in undesirable side-effects for the network path, such as route or DNS flapping. A DOTS server therefore MAY continue mitigating for a mutually negotiated period after receiving the DOTS client’s request to stop.

A server MAY refuse to engage in coordinated attack response with a client. To make the status of a client’s request clear, the server MUST indicate in server signals whether client-initiated mitigation is active. When a client-initiated mitigation is active, and threat handling details such as mitigation scope and statistics are available to the server, the server SHOULD include those details in server signals sent to the client. DOTS clients SHOULD take mitigation statistics into account when deciding whether to request the DOTS server cease mitigation.

OP-006 Mitigation Scope: DOTS clients MUST indicate the desired address space coverage of any mitigation, for example by using Classless Internet Domain Routing (CIDR) [RFC1518],[RFC1519] prefixes, [RFC2373] for IPv6 prefixes, the length/prefix convention established in the Border Gateway Protocol (BGP) [RFC4271], or by a prefix group alias agreed upon with the server through the data channel. If there is additional information available narrowing the scope of any requested attack response, such as targeted port range, protocol, or service, clients SHOULD include that information in client signals.

As an active attack evolves, clients MUST be able to adjust as necessary the scope of requested mitigation by refining the address space requiring intervention.

2.3. Data channel requirements

The data channel is intended to be used for bulk data exchanges between DOTS agents. Unlike the signal channel, which must operate nominally even when confronted with despite signal degradation due to packet loss, the data channel is not expected to be constructed to deal with attack conditions. As the primary function of the data channel is data exchange, a reliable transport is required in order for DOTS agents to detect data delivery success or failure.

The data channel should be adaptable and extensible. We anticipate the data channel will be used for such purposes as configuration or resource discovery. For example, a DOTS client may submit to the DOTS server a collection of prefixes it wants to refer to by alias when requesting mitigation, to which the server would respond with a success status and the new prefix group alias, or an error status and message in the event the DOTS client’s data channel request failed. The transactional nature of such data exchanges suggests a separate set of requirements for the data channel, while the potentially sensitive content sent between DOTS agents requires extra precautions to ensure data privacy and authenticity.

DATA-001 Reliable transport: Transmissions over the data channel may be transactional, requiring reliable, in-order packet delivery.

DATA-002 Data privacy and integrity: Transmissions over the data channel may contain sensitive information or instructions from the remote DOTS agent. Theft or modification of data channel transmissions could lead to information leaks or malicious transactions on behalf of the sending agent. (See Security Considerations below.) Consequently data sent over the data channel MUST be encrypted and authenticated using current industry best practices.

DATA-003 Mutual authentication: DOTS agents MUST mutually authenticate each other before data may be exchanged over the data channel. DOTS agents MAY take additional steps to authorize data exchange, as in the prefix group example above, before accepting data over the data channel. The form of authentication and authorization is unspecified.

DATA-004 Black- and whitelist management: DOTS servers SHOULD provide methods for DOTS clients to manage black- and white-lists of source addresses of traffic destined for addresses belonging to a client.

For example, a DOTS client should be able to create a black- or whitelist entry; retrieve a list of current entries from either list; update the content of either list; and delete entries as necessary.

How the DOTS server determines client ownership of address space is not in scope.

2.4. Data model requirements

TODO

3. Congestion Control Considerations

The DOTS signal channel will not contribute measurably to link congestion, as the protocol's transmission rate will be negligible regardless of network conditions. Bulk data transfers are performed over the data channel, which should use a reliable transport with built-in congestion control mechanisms, such as TCP.

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.

5. Change Log

5.1. 00 revision

2015-10-15

5.2. Initial revision

2015-09-24 Andrew Mortensen

6. References

6.1. Normative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <http://www.rfc-editor.org/info/rfc768>.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <http://www.rfc-editor.org/info/rfc793>.

[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>.

6.2. Informative References

[RFC1518] Rekhter, Y. and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, DOI 10.17487/RFC1518, September 1993, <http://www.rfc-editor.org/info/rfc1518>.

[RFC1519] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, DOI 10.17487/RFC1519, September 1993, <http://www.rfc-editor.org/info/rfc1519>.

[RFC2373] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, DOI 10.17487/RFC2373, July 1998, <http://www.rfc-editor.org/info/rfc2373>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <http://www.rfc-editor.org/info/rfc4271>.

Authors' Addresses

Andrew Mortensen
Arbor Networks, Inc.
2727 S. State St
Ann Arbor, MI 48104
United States

Email: amortensen@arbor.net

Robert Moskowitz
HTT Consulting
Oak Park, MI 42837
United States

Email: rgm@htt-consult.com

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tireddy@cisco.com

DOTS WG
Internet-Draft
Intended status: Informational
Expires: April 21, 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

October 19, 2015

Use cases for DDoS Open Threat Signaling
draft-ietf-dots-use-cases-00.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.

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 April 21, 2016.

Table of Contents

1. Requirements notation

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

2. 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.

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 and/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.

3. Terminology and Acronyms

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

- DDoS: A distributed denial-of-service attack. DDoS attacks are intended to cause a negative impact on the availability of servers, services, applications, and/or other functionality of an attack target.

- Attack target: The intended target of a DDoS attack.

- Attack telemetry: Collected network traffic characteristics enabling the detection, classification, and in many cases traceback of a DDoS attack.

- Mitigation: A defensive response against a detected DDoS attack, performed by an entity in the network path between attack sources and the attack target, either through inline deployment or some form of traffic diversion, consisting of one or more countermeasures. The form a given mitigation takes is out of scope for this document.

- Countermeasure: An action or set of actions taken by a mitigator to evaluate and filter out a significant proportion of DDoS attack traffic while forwarding onwards a significant proportion of legitimate traffic directed towards an attack target.

4. 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, then 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 which have been implemented in a largely proprietary manner by various DDoS mitigation solution vendor and/or service providers, resulting in vendor/service lock-in and mutually

incompatible solutions/services. The overarching goal of the DOTS effort is to provide a standards-based mechanism to allow heterogeneous DDoS mitigation solutions and services to be woven together in order to allow broader, more pervasive adoption of coordinated DDoS defense.

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.

Another important goal of DOTS is interoperability and coordination via a common standards-based mechanism between multiple DDoS mitigation service providers contemporaneously engaged in defending the same organization against DDoS attacks. 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.

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

4.1. Primary Use Cases

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

In this scenario, 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.

All DOTS messages exchanged between the DOTS clients and DOTS servers in this use case may be communicated directly between the 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.

(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. 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 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 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(j) The CPE or PE DDoS mitigators transmit a DOTS mitigation service termination request to the DOTS servers. [The scope, format, and content of these messages must be codified by the DOTS WG.] 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(m) The CPE or PE DDoS mitigators transmit a DOTS mitigation termination status acknowledgement to the DOTS servers. [The scope, format, and content of these messages must be codified by the DOTS WG.]

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

In this scenario, 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.

All DOTS messages exchanged between the DOTS clients and DOTS servers in this use case may be communicated directly between the 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.

(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 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the requesting requesting network infrastructure elements. The scope, format, and content of these messages must be codified by the DOTS WG.

(f) While DDoS mitigation services are active, the network infrastructure elements may optionally regularly transmit DOTS

mitigation efficacy updates to the relevant DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(i) The network infrastructure elements transmit a DOTS mitigation service termination request to the DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG. 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(l) The network infrastructure elements transmit a DOTS mitigation termination status acknowledgement to the DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG.

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

In this scenario, 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.

All DOTS messages exchanged between the DOTS clients and DOTS servers in this use case may be communicated directly between the 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.

(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. [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 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 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(i) The CPE or PE Attack Telemetry Detection/Classification Systems transmit a DOTS mitigation service termination request to the DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG. 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(l) The CPE or PE Attack Telemetry Detection/Classification Systems transmit a DOTS mitigation termination status acknowledgement to the DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG.

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

In this scenario, 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.

All DOTS messages exchanged between the DOTS clients and DOTS servers in this use case may be communicated directly between the 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.

(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. 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 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. [The scope, format, and content of these messages must be codified by the DOTS WG.

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the requesting services or applications. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(i) The targeted services or applications transmit a DOTS mitigation service termination request to the DOTS servers. [The scope, format, and content of these messages must be codified by the DOTS WG.] 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. The scope, format, and content of these messages must be codified by the DOTS WG.

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

scope, format, and content of these messages must be codified by the DOTS WG.

4.1.5. Successful Manual Web Portal Request to Upstream Mitigator

In this scenario, 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.

All DOTS messages exchanged between the DOTS client and DOTS servers in this use case may be communicated directly between the DOTS client 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.

(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. [The scope, format, and content of these messages must be codified by the DOTS WG.] 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. [The scope, format, and content of these messages must be codified by the DOTS WG.

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the Web portal. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. [The scope, format, and content of these messages must be codified by the DOTS WG.

(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). The scope, format, and content of these messages must be codified by the DOTS WG.

(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 Web portal indicating that DDoS mitigation services have been terminated. The scope, format, and content of these messages must be codified by the DOTS WG.

(l) The Web portal transmits a DOTS mitigation termination status acknowledgement to the DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG.

4.1.6. Successful Manual Mobile Device Application Request to Upstream Mitigator

In this scenario, 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.

All DOTS messages exchanged between the DOTS client and DOTS servers in this use case may be communicated directly between the DOTS client 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.

(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. [The scope, format, and content of these messages must be codified by the DOTS WG.] 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(e) While DDoS mitigation services are active, the DOTS servers regularly transmit DOTS mitigation status updates to the mobile application. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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). The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(l) The mobile application transmits a DOTS mitigation termination status acknowledgement to the DOTS servers. The scope, format, and content of these messages must be codified by the DOTS WG.

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

In this scenario, 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.

All DOTS messages exchanged between the DOTS clients and DOTS servers in this use case may be communicated directly between the 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.

(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. 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 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 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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG.

(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. The scope, format, and content of these messages must be codified by the DOTS WG. 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.

4.2. Ancillary Use Cases

4.2.1. 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.

4.2.2. 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.

4.2.3. 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.

5. 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.

6. IANA Considerations

7. Acknowledgments

8. References

8.1. Normative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <http://www.rfc-editor.org/info/rfc768>.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <http://www.rfc-editor.org/info/rfc793>.

[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>.

[RFC1518] Rekhter, Y. and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, DOI 10.17487/RFC1518, September 1993, <http://www.rfc-editor.org/info/rfc1518>.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <http://www.rfc-editor.org/info/rfc4632>.

8.2. Informative References

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

[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

Stephane 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
US

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

DOTS
Internet-Draft
Intended status: Informational
Expires: April 21, 2016

K. Nishizuka
NTT Communications

October 19, 2015

Inter-Domain DOTS Use Cases
draft-nishizuka-dots-inter-domain-usecases-00

Abstract

This document describes inter-domain use cases of the DDoS Open Threat Signaling(DOTS).

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 April 22, 2016.

Table of Contents

1. Introduction

Maximum size of DDoS attack is increasing. According to a report from Cloudflare[Cloudflare], in 2013, over 300 Gbps DDoS attack against Spamhaus was observed which exploited DNS reflection mechanism to create massive attack with intention to overwhelm the capacity of the targeted system.

If this trend continued, the volume of DDoS attack will exceed preparable anti-DDoS capability by one organization mostly in the aspect of cost. Moreover, possibility of DDoS attack is unpredictable, so it is not realistic that every organization prepare sufficient anti-DDoS system.

This problem could be solved by sharing anti-DDoS system over multi-organizations. We can share the burden of protection against DDoS attack by inter-domain cooperation. To accomplish this, we need a framework which use common interface to call for protection.

To describe the mechanism of such a framework, we classified inter-domain use cases into three models.

1. Multi-home Model (one supplicant and multi mitigators)

2. Cloud Model (multi supplicants and one mitigator)

3. Delegation Model (both sides of supplicant and mitigator)

By blocking DDoS attack with inter-domain cooperation, average usage of DDoS mitigation equipment will increase. This will leverage total capacity of anti-DDoS system in all over the internet. With this mechanism, we can manage DDoS attacks which exceed the capacity of its own platform.

At the same time, it might be needed to convey information of amount of processed threat traffic which would be used to charge other organization each other. However this kind of information is out of scope of DOTS.

2. Terminology

supplicant: call for an anti-DDoS action to a mitigator. It could be a service under attack itself. Also, it could be a monitoring system which inspect the traffic towards the service by netflow/sflow or DPI, from which it can detect DDoS attack in the traffic. The minimum requirement to supplicant is that it must know which IP address is under attack and convey it to a mitigator by DOTS protocol.

mitigator: protect a service from DDoS attack. It can use blackholing, ACLs, flowspec, rate-limit, dedicated DDoS mitigation devices and other methods depending on its capabilities. It must be preprovisioned to determine a DDoS protection entity. It starts DDoS protection based on information provided by a supplicant. The minimum information is IP address of the service which it must protect from the DDoS attack. Other information like source IP address, port, type of DDoS, etc. provided by the supplicant are optional. The optional information may be used, but it might be overridden by the mitigator according to the on-going attack.

Other terminology and acronyms are inherited from [I-D.draft-mglt-dots-use-cases]

3. DDoS Protection Scenario

DDoS protection can be divided into two stages.

o Provisioning stage

Before getting attacked by malicious traffic, a supplicant needs capacity building with a mitigator in advance. In this provisioning stage, following information should be provided to the mitigator side to prepare for DDoS attack:

1. Protection capability

2. Restriction on the range of IP addresses and ports

3. Return path information of the mitigated traffic

4. Authorization information to restrict the supplicant

These informations can be conveyed off the wire, thus this is out of scope of DOTS. However, provisioning stage is very important to protect the service, therefore we describes how DDoS protection works comprehensively.

o Signaling stage

After getting attacked, we need to signal SOS information immediately if the service has not implemented any other anti-DDoS system except preprovisioned DDoS mitigation. In this signaling stage, the supplicant signals targeted IP address to the mitigator with authorization information. The mitigator decides to protect the system based on the preprovisioned information. This signaling should have characteristics as follows:

1. Common transport and schema

2. Secure signaling

Even in the signaling stage, preprovisioned information can be changed according to the DDoS attack vector. However, provisioning and signaling must be separated to keep DOTS requirements simple.

3.1. Provisioning Stage

In this section, we describe how preprovisioned information is used to protect a service. In the provisioning stage, before getting attacked, the operator of the service register following informations to a mitigator to protect the service correctly and effectively.

3.1.1. Protection Capability

Protection capability is consist of three informations: protection method, protection threshold and traffic capacity.

o protection method

Available protection methods of mitigator may be selectable, which include blackholing, ACLs, flowspec, dedicated DDoS appliances, etc. These methods have their own max capacity. Therefore, protection threshold should be determined in advance according to the traffic capacity of the method.

In the case of blackholing, it stops the traffic destined to the service totally. In a way, the "denial" of service is successful except in the case of selective blackhole. However, the capacity of the blackholing is rather higher than other methods because it just divert traffic to null0 interface of routers.

On the other hand, in the case of DDoS mitigation appliances, only the malicious traffic will be discarded on the box and the scrubbed normal traffic will be returned to the original service thus service continuity will be kept, though there is possibility of false positives and false negatives. However, the total volume of processable traffic is limited to the capacity of the hardware. To reduce the possibility of the mis-classification, which type of DDoS attack will be processed and which countermeasures will be applied to should be determined in the provisioning stage.

o protection threshold

Protection threshold defines when the appropriate method should be invoked to start protection. Typical threshold is traffic volume(bps/pps) of the attack. Depending on the type of the service, the appropriate threshold differs. If the threshold is not appropriate, possibility of false positives and false negatives increases. For example, if the service is widely used content server, low threshold of SYN attack protection(rate-limit) could cause failure of normal transaction.

o traffic capacity

Traffic capacity is protectable total volume[bps/pps] of DDoS traffic which include both malicious traffic and normal traffic. This capacity should be negotiated carefully because it could affect the service directly. From the point of view of the mitigator, maximum duration and number of protection could be limited to protect the DDoS mitigation system from exclusive occupancy.

If the protection capability of one mitigator is insufficient to a service, DOTS can provide capacity leverage to both the service and the mitigator.

3.1.2. Restriction on the Range of IP Addresses and Ports

In the provisioning stage, the service should register the range of IP addresses which they need to protect to the mitigator. Without this restriction, they can use anti-DDoS system to protect any other organization. Especially, in case of blackholing, they can abuse the system by blocking all of the traffic to the other organization.

In addition, they can register range of source IP address/port and destination IP address/port as a whitelist. If they know some range of 5 tuples which never include DDoS traffic, they can exclude it from the target of anti-DDoS protection, which reduce the possibility of false positive.

3.1.3. Return Path Information of the Mitigated Traffic

In many cases, DDoS mitigator controls traffic to divert DDoS attack traffic to its own domain to deal with it. It classifies the traffic into malicious traffic and normal traffic. Normal traffic should be returned to the original server, however simply returning traffic to the internet can cause routing loop because the returning traffic could re-enter the diversion path again. To avoid this routing loop, the returning path should be provisioned in advance. If there is no dedicated line between the mitigator and the service, tunnel technology such as GRE[RFC2784] can be used. In that case, tunnel information should be preprovisioned. In general, next-hop and prefix information should be provided to the mitigator to determine the returning path of the mitigated traffic.

3.1.4. Authorization Information to Restrict the Supplicant

After the provisioning, the mitigator should limit the usage of the provisioned DDoS protection entity to the legitimate supplicant. Only authorized supplicant can trigger the anti-DDoS action. If the supplicant was not restricted, a spoofed signal could abuse the mitigator. Also, the system should be protected from replay attack.

3.2. Signaling Stage

After the provisioning stage, the authorization information of the DDoS protection entity will be supplied to a supplicant. Then, the supplicant can call for help to the DDoS mitigator by signaling mandatory information.

3.2.1. Signaling Information

The mandatory information which should be included in the signaling is as follows:

- IP address of defence target
- Instruction (Start/Stop)
- Authorization information

Suppose a supplicant, which is the service itself or monitoring system, can know that the service is under a severe DDoS attack. After the detecting the DDoS attack, the supplicant records attacked IP address(es). Adding the authorization information provided in advance, it signals protection-start-instruction packet to the mitigator including IP address of defence target.

The mitigator which recieved the signal reacts to start mitigation. First, it checks the authorization information to decide the signaling is legitimate or not. If failed, it never react. If succeeded, it checks IP address with according DDoS protection entity. Second, If the IP address was included in the range which was declaired in advance, it starts mitigation. The protection method will be selected appropoately according to the provisioned protection capability. Finally, it classifies malicious traffic and normal traffic, then return the normal traffic to the service in specified returning path.

The supplicant can stop the mitigation by sending protection-stop-instruction packet. However, in some case, it is difficult to know whether the DDoS attack has ended or not from the monitoring point of the supplicant.

The following informations are useful for mitigators in many cases but they are optional.

- Attack ID
- (Average/Maximum/Currrent)Traffic volume[bps/pps]
- Severity
- Type of attack
- Protection method
- Src IP/Port

- Dst Port

- Attack start time

We describe the reason why these informations are not mandatory.

- Attack ID

Attack ID could be assigned by a supplicant. By recieving the attack ID, a mitigator can tell the attack vector is the same or not from the observation of the supplicant. However, regardless of the provided attack ID, the behavior of DDoS protection will not change. Therefore this is optional information.

- (Average/Maximum/Currrent)Traffic volume[bps/pps]

Traffic volume information can be used to determine protection method. However, in the case of massive DDoS attack, the circuit connected to the internet from the service could be saturated by the traffic, so there is no way to know how much traffic is incoming on the saturated link. Thus, traffic volume information provided by the supplicant is unreliable. That is why this is optional information.

- Severity

Severity information can be used to determine protection method. However, in many cases, DDoS attack vectors change time to time, so there is no constant index of severity. Moreover, the monitoring system on the service side can look through the important attack vector which is very severe to the service, so the severity must be overwritten by the mitigator if it can inspect the traffic more deeply. Therefore this is optional information.

- Type of attack

Similar to severity information, type of attack declared by the monitoring system on the service side is unreliable. Decision of the type of attack must be overwritten by the mitigator if it can inspect the traffic more deeply. Therefore this is optional information.

- Protection method

The supplicant can convey preferable protection method information, which could be used to change the behavior of the mitigator. However, depending on the usage situation, the mitigator could override the protection method. Therefore this is optional information.

- Src IP/Port

In some cases, source IP/Port of the DDoS attack are spoofed. They widely vary and continue changing. Thus, the mitigator can not depend on the Src IP/Port information from the supplicant. Therefore this is optional information.

- Dst Port

Destination port of the DDoS attack can be changed by the attacker if they observed the attack on the port is not effective. Similar to Src IP/Port information, this is optional information.

- Attack start time

Attack start time information can indicate the severity of the attack. The mitigator can find the attack effectively by that inforamtion if it has a constant monitoring system. However, this is optional information.

3.2.2. Common Transport and Schema

To convey the information listed in the previous section, DOTS WG will define a common transport and schema. These are under discussion on Mailing List based on the draft [I-D.draft-reddy-dots-transport]. Defining these common transport and schema is out of scope of this draft. We note that, with a common transport and schema, we can share the burden of protection against DDoS attack in inter-domain model, which is described in Section.4.

3.2.3. Secure Signaling

Secure signaling is fundamental requirement to the DOTS signaling protocol. Only the legitimate supplicants can use the mitigator. Restriction can be accomplished by existing authentication and authorization methodologies. Signaling must be encrypted to avoid man-in-the-middle attack. To deal with the unreliable transport on the link under attack, signaling should have idempotency. Also authorization information must be securely exchanged in the provisioning stage. Though these characteristics are important, defining the signaling method is out of scope of this draft.

3.3. After DDoS Protection

After the DDoS protection was kicked by signaling, some information derived from the mitigator is useful to the operators of the service.

- Status of ongoing protection

Status of the protection(The attack is ongoing or not) will be used to determine that the system is already safe without the protection. The mitigator should have interface from which the supplicant or the operator of the service can get the status of the protection.

o Attack information

The operator of the service will eager to know what kind of attack was pointed to the service. Then, they can study how to try to find the best plan to cope with the situation.

o Number of the dropped packets

Number of the dropped packets can be used to create the billing data. Some DDoS mitigator may have data quantity charging system to account the supplicant based on the usage of their resources.

How to convey these information is indispensable issue of inter-domain DDoS protection. However, we note that these are out of scope of DOTS.

4. Inter-Domain Dots Use Cases

We classified inter-domain use cases into three models. In these models, the signaling packets traverse over multi domains. They utilize the common interface to the DDoS protection entities which are located in the multiple domains. We assume that the provisioning stage has finished in all mitigators, so by sending signaling packets, the mitigators start the according protections and return scrubbed traffic to the service in specified return path.

4.1. Usecase 1: Multi-home Model

In the multi-home model, there are one supplicant and multi mitigators. The supplicant can use both mitigators.

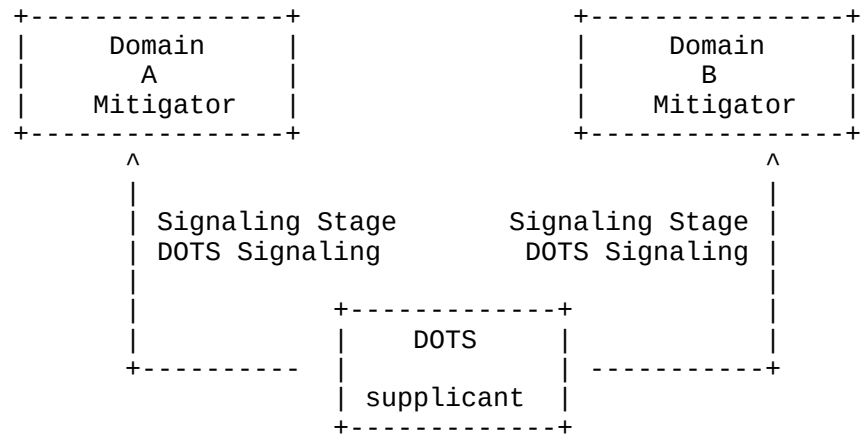

Figure 1: Usecase 1: Multi-home Model

An example of this situation is that a content provider is connected to two transit providers. When the content provider get attacked, the DDoS traffic will come from transit A and B. Signaling to the mitigator in transit A can stop only the DDoS traffic from transit A, and vice verse. Though the provision method will be different, the signaling interfaces are common if the both mitigators are using dots framework. After detecting the DDoS attack, the supplicant will send the signaling packet to the both mitigators at the same time. Common interface of DOTS signaling will shorten the lead time of the DDoS protection on both transits.

4.2. Usecase 2: Cloud Model

In the cloud model, there are multi supplicants and one mitigator. The mitigator accepts signals from multi supplicants in multiple domains.

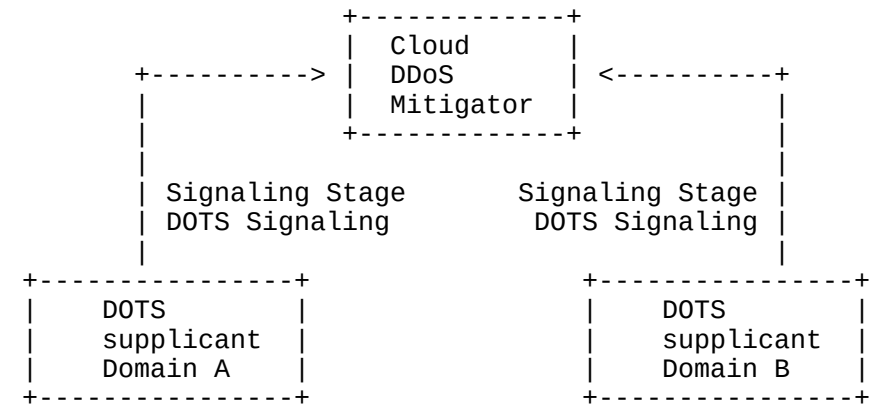

Figure 2:Usecase 2: Cloud Model"

An example of this situation is cloud type of DDoS mitigation service provider. Cloud type of DDoS mitigation service providers divert traffic to its own domain using routing protocols, that is BGP route injection. Though they need to provision the returning path mostly on the tunnel interface because they are not directly connected to the domains of the supplicants, they can accomodate multiple domains remotely.

4.3. Usecase 3: Delegation Model

In the delegation model, a mitigator has both sides of supplicant and mitigator.

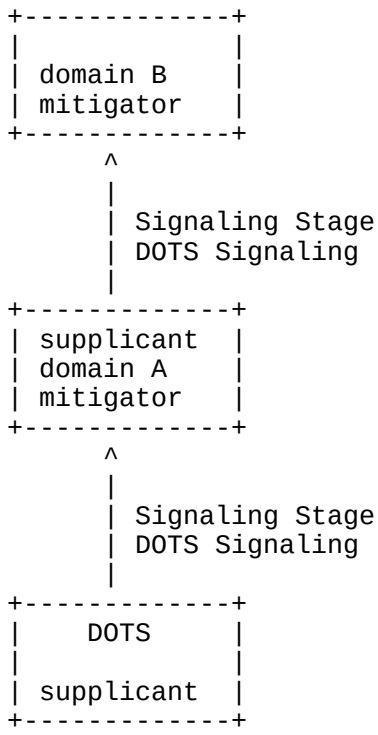

Figure 3: Usecase 3: Delegation Model

If the capacity of the mitigator is insufficient in comparison with ongoing DDoS attack, the mitigator can be a supplicant which call for protection in other domain. The provisioning of the mitigator in domain B can be done by the mitigator in domain A as a supplicant in advance. By just relaying the DOTS signaling information to the mitigator in domain B, the mitigator in domain A can utilize DDoS protection of doamin B. The original supplicant might not notice that the mitigation was delegated to other domain. Even if the capacity is sufficient, in some cases, it is effective to delegate the protection to upstream domain. Stopping DDoS traffic at an ingress border will reduce unnecessary forwarding. The mitigator can delegate the burden of the mitigation, therefore they can accomodate more services which exceed the capacity of its own platform.

A mitigator can be a broker which select appropriate DDoS mitigators according to the capacities and the field of expertise of the mitigators. In this case, billing data could be more important to adjust the cost distribution fairly.

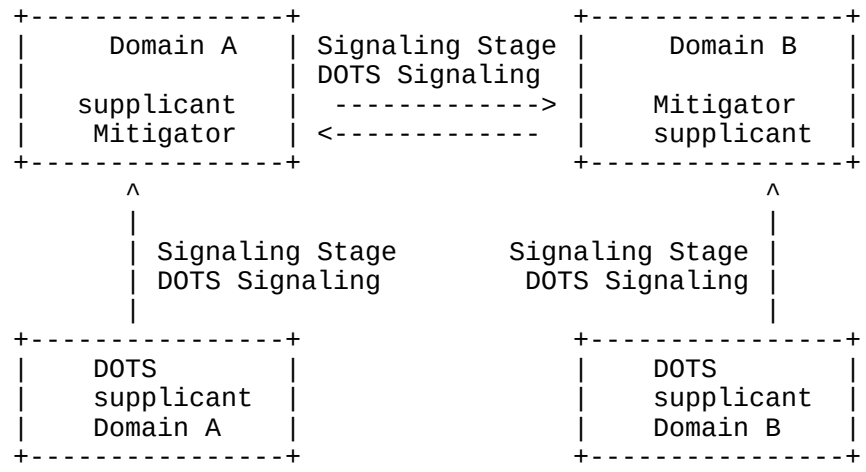

Figure 4: Cooperative DDoS Mitigation with DOTS Signaling

The figure.4 describes a minor changed version of the delegation model. The supplicants and mitigators can signal each other with DOTS signaling. They can ask for help each other. In this model, we can leverage total capacity of anti-DDoS system in all over the internet.

5. Security Considerations

As described in Section.3.2.3, secure signaling is fundamental requirement to the DOTS signaling protocol. Only the legitimate supplicants can use the mitigator. Authorization information must be securely exchanged in the provisioning stage.

6. IANA Considerations

No need to describe any request regarding number assignment.

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>.

[RFC2784] D. Farinacci., T. Li., S. Hanks., D. Meyer., and P. Traina., "Generic Routing Encapsulation (GRE), March 2000".

[I-D.draft-mglt-dots-use-cases] D. Migault, Ed., "DDoS Open Threat Signaling use cases, draft-mglt-dots-use-cases-00 (work in progress), April 2015".

[I-D.draft-reddy-dots-transport] T. Reddy., D. Wing., P. Patil., M. Geller., M. Boucadair., and R. Moskowitz., "Co-operative DDoS Mitigation, October 2015".

7.2. URL References

[Cloudflare] Cloudflare, "https://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho/".

Author's Address

Kaname Nishizuka
NTT Communications
GranPark 16F
3-4-1 Shibaura, Minato-ku, Tokyo
108-8118,Japan

EMail: kaname@nttv6.jp

DOTS
Internet-Draft
Intended status: Standards Track
Expires: April 20, 2016

T. Reddy
D. Wing
P. Patil
M. Geller
Cisco
M. Boucadair
France Telecom
R. Moskowitz
HTT Consulting

October 18, 2015

Co-operative DDoS Mitigation
draft-reddy-dots-transport-01

Abstract

This document discusses mechanisms that a DOTS client can use, when it detects a potential Distributed Denial-of-Service (DDoS) attack, to signal that the DOTS client is under an attack or request an upstream DOTS server to perform inbound filtering in its ingress routers for traffic that the DOTS client wishes to drop. The DOTS server can then undertake appropriate actions (including, blackhole, drop, rate-limit, or add to watch list) on the suspect traffic to the DOTS client, thus reducing the effectiveness of the attack.

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 April 20, 2016.

Table of Contents

1. Introduction

A distributed denial-of-service (DDoS) attack is an attempt to make machines or network resources unavailable to their intended users. In most cases, sufficient scale can be achieved by compromising enough end-hosts and using those infected hosts to perpetrate and amplify the attack. The victim in this attack can be an application server, a client, a router, a firewall, or an entire network, etc. The reader may refer, for example, to [REPORT] that reports the following:

- Very large DDoS attacks above the 100 Gbps threshold are experienced.

- DDoS attacks against customers remain the number one operational threat for service providers, with DDoS attacks against infrastructures being the top concern for 2014.

- Over 60% of service providers are seeing increased demand for DDoS detection and mitigation services from their customers (2014), with just over one-third seeing the same demand as in 2013.

In a lot of cases, it may not be possible for an enterprise to determine the cause for an attack, but instead just realize that certian resources seem to be under attack. The document proposes that, in such cases, the DOTS client just inform the DOTS server that the enterprise is under a potential attack and that the DOTS server monitor traffic to the enterprise to mitigate any possible attack. This document also describes a means for an enterprise, which act as DOTS clients, to dynamically inform its DOTS server of the IP addresses or prefixes that are causing DDoS. A DOTS server can use this information to discard flows from such IP addresses reaching the customer network.

The proposed mechanism can also be used between applications from various vendors that are deployed within the same network, some of them are responsible for monitoring and detecting attacks while others are responsible for enforcing policies on appropriate network elements. This cooperations contributes to a ensure a highly automated network that is also robust, reliable and secure. The advantage of the proposed mechanism is that the DOTS server can provide protection to the DOTS client from bandwidth-saturating DDoS traffic.

How a DOTS server determines which network elements should be modified to install appropriate filtering rules is out of scope. A variety of mechanisms and protocols (including NETCONF) may be considered to exchange information through a communication interface between the server and these underlying elements; the selection of appropriate mechanisms and protocols to be invoked for that interfaces is deployment-specific.

Terminology and protocol requirements for co-operative DDoS mitigation are obtained from [I-D.mortensen-dots-requirements].

2. Notational Conventions

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

3. Solution Overview

Network applications have finite resources like CPU cycles, number of processes or threads they can create and use, maximum number of simultaneous connections it can handle, limited resources of the

control plane, etc. When processing network traffic, such an application uses these resources to offer its intended task in the most efficient fashion. However, an attacker may be able to prevent the application from performing its intended task by causing the application to exhaust the finite supply of a specific resource.

TCP DDoS SYN-flood is a memory-exhaustion attack on the victim and ACK-flood is a CPU exhaustion attack on the victim. Attacks on the link are carried out by sending enough traffic such that the link becomes excessively congested, and legitimate traffic suffers high packet loss. Stateful firewalls can also be attacked by sending traffic that causes the firewall to hold excessive state and the firewall runs out of memory, and can no longer instantiate the state required to pass legitimate flows. Other possible DDoS attacks are discussed in [RFC4732].

In each of the cases described above, if a network resource detects a potential DDoS attack from a set of IP addresses, the network resource (DOTS client) informs its servicing router (DOTS relay) of all suspect IP addresses that need to be blocked or black-listed for further investigation. DOTS client could also specify protocols and ports in the black-list rule. That DOTS relay in-turn propagates the black-listed IP addresses to the DOTS server and the DOTS server blocks traffic from these IP addresses to the DOTS client thus reducing the effectiveness of the attack. The DOTS client periodically queries the DOTS server to check the counters mitigating the attack. If the DOTS client receives response that the counters have not incremented then it can instruct the black-list rules to be removed. If a blacklisted IPv4 address is shared by multiple subscribers then the side effect of applying the black-list rule will be that traffic from non-attackers will also be blocked by the access network.

If a DOTS client cannot determine the IP address(s) that are causing the attack, but is under an attack nonetheless, the DOTS client can just notifiy the DOTS server that it is under a potential attack and request that the DOTS server take precautionary measures to mitigate the attack.

4. Protocol for Signal Channel: HTTP REST

A DOTS client can use RESTful APIs discussed in this section to signal/inform a DOTS server of an attack or any desired IP filtering rules.

4.1. SOS

The following APIs define the means to signal an SOS from a DOTS client to a DOTS server.

TBD: SOS messages SHOULD be exchanged over DTLS over UDP.

4.1.1. Signal SOS

An HTTP POST request will be used to signal SOS to the DOTS server.

```
POST {scheme}://{host}:{port}/.well-known/{version}/{URI suffix for SOS}
Accept: application/json
Content-type: application/json
{
   "policy-id": number,
   "target-ip": string,
   "target-port": string,
   "target-protocol": string,
}
```

Figure 1: POST to signal SOS

The header fields are described below.

policy-id: Identifier of the policy represented using a number. This identifier must be unique for each policy bound to the DOTS client. This identifier must be generated by the client and used as an opaque value by the server. This document does not make any assumption about how this identifier is generated.

target-ip: A list of addresses or prefixes under attack. This is an optional attribute.

target-port: A list of ports under attack. This is an optional attribute.

target-protocol: A list of protocols under attack. Valid protocol values include tcp, udp, sctp and dccp. This is an optional attribute.

Note: administrative-related clauses may be included as part of the request (such a contract Identifier or a customer identifier). Those clauses are out of scope of this document.

To avoid SOS message fragmentation and the consequently decreased probability of message delivery, DOTS agents MUST ensure that the DTLS record MUST fit within a single datagram. DOTS agents can

exploit the fact that the IP specification [RFC0791] specifies that IP packets up to 576 bytes should never need to be fragmented, thus sending a maximum of 500 bytes of SOS message over a UDP datagram will generally avoid IP fragmentation.

The following example shows POST request to signal that a Web-Service is under attack.

```
POST https://www.example.com/.well-known/v1/SOS
Accept: application/json
Content-type: application/json
 {
   "policy-id": 123321333242,
   "target-ip": "2002:db8:6401::1",
   "target-port": "443",
   "target-protocol": "tcp",
 }
```

Figure 2: POST to signal SOS

4.1.2. Recall SOS

An HTTP DELETE request will be used to delete an SOS signaled to the DOTS server.

```
DELETE {scheme}://{host}:{port}/.well-known/{URI suffix for SOS}
Accept: application/json
Content-type: application/json
 {
   "policy-id": number
 }
```

Figure 3: Recall SOS

4.1.3. Retrieving SOS

An HTTP GET request will be used to retrieve an SOS signaled to the DOTS server.

1) To retrieve all SOS signaled by the DOTS client.

```
GET {scheme}://{host}:{port}/.well-known/{URI suffix for SOS}
```

2) To retrieve a specific SOS signaled by the DOTS client.

```
GET {scheme}://{host}:{port}/.well-known/{URI suffix for SOS}
Accept: application/json
Content-type: application/json
 {
   "policy-id": number
 }
```

Figure 4: GET to retrieve the rules

4.2. REST

A DOTS client could use HTTP to provision and manage filters on the DOTS server. The DOTS client authenticates itself to the DOTS relay, which in turn authenticates itself to a DOTS server, creating a two-link chain of transitive authentication between the DOTS client and the DOTS server. The DOTS relay validates if the DOTS client is authorized to signal the black-list rules. Likewise, the DOTS server validates if the DOTS relay is authorized to signal the black-list rules. To create or purge filters, the DOTS client sends HTTP requests to the DOTS relay. The DOTS relay acts as an HTTP proxy, validates the rules and proxies the HTTP requests containing the black-listed IP addresses to the DOTS server. When the DOTS relay receives the associated HTTP response from the HTTP server, it propagates the response back to the DOTS client.

If an attack is detected by the DOTS relay then it can act as a HTTP client and signal the black-list rules to the DOTS server. Thus the DOTS relay plays the role of both HTTP client and HTTP proxy.

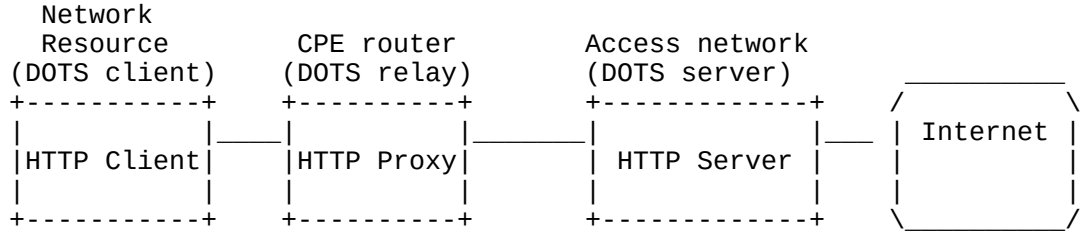

Figure 5

JSON [RFC7159] payloads can be used to convey both filtering rules as well as protocol-specific payload messages that convey request parameters and response information such as errors.

The figure above explains the protocol with a DOTS relay. The protocol is equally applicable to scenarios where a DOTS client directly talks to the DOTS server.

4.2.1. Filtering Rules

The following APIs define means for a DOTS client to configure filtering rules on a DOTS server.

4.2.1.1. Install filtering rules

An HTTP POST request will be used to push filtering rules to the DOTS server.

```
POST {scheme}://{host}:{port}/.well-known/{version}/{URI suffix for filtering}
Accept: application/json
Content-type: application/json
{
   "policy-id": number,
   "traffic-protocol": string,
   "source-protocol-port": string,
   "destination-protocol-port": string,
   "destination-ip": string,
   "source-ip": string,
   "lifetime": number,
   "traffic-rate" : number,
 }
```

Figure 6: POST to install filtering rules

The header fields are described below.

policy-id: Identifier of the policy represented using a number. This identifier must be unique for each policy bound to the same downstream network. This identifier must be generated by the client and used as an opaque value by the server. This document does not make any assumption about how this identifier is generated.

traffic-protocol: Valid protocol values include tcp and udp.

source-protocol-port: For TCP or UDP or SCTP or DCCP: the source range of ports (e.g., 1024-65535).

destination-protocol-port: For TCP or UDP or SCTP or DCCP: the destination range of ports (e.g., 443-443). This information is useful to avoid disturbing a group of customers when address sharing is in use [RFC6269].

destination-ip: The destination IP addresses or prefixes.

source-ip: The source IP addresses or prefixes.

lifetime: Lifetime of the policy in seconds. Indicates the validity of a rule. Upon the expiry of this lifetime, and if the request is not reiterated, the rule will be withdrawn at the upstream network. A null value is not allowed.

traffic-rate: This field carries the rate information in IEEE floating point [IEEE.754.1985] format, units being bytes per second. A traffic-rate of '0' should result on all traffic for the particular flow to be discarded.

The relative order of two rules is determined by comparing their respective policy identifiers. The rule with lower numeric policy identifier value has higher precedence (and thus will match before) than the rule with higher numeric policy identifier value.

Note: administrative-related clauses may be included as part of the request (such a contract Identifier or a customer identifier). Those clauses are out of scope of this document.

The following example shows POST request to block traffic from attacker IPv6 prefix 2001:db8:abcd:3f01::/64 to network resource using IPv6 address 2002:db8:6401::1 to provide HTTPS web service.

```
POST https://www.example.com/.well-known/v1/filter
Accept: application/json
Content-type: application/json
 {
   "policy-id": 123321333242,
   "traffic-protocol": "tcp",
   "source-protocol-port": "1-65535",
   "destination-protocol-port": "443",
   "destination-ip": "2001:db8:abcd:3f01::/64",
   "source-ip": "2002:db8:6401::1",
   "lifetime": 1800,
   "traffic-rate": 0,
 }
```

Figure 7: POST to install black-list rules

4.2.1.2. Remove filtering rules

An HTTP DELETE request will be used to delete filtering rules programmed on the DOTS server.

```
DELETE {scheme}://{host}:{port}/.well-known/{URI suffix for filtering}
Accept: application/json
Content-type: application/json
 {
   "policy-id": number
 }
```

Figure 8: DELETE to remove the rules

4.2.1.3. Retrieving installed filtering rules

An HTTP GET request will be used to retrieve filtering rules programmed on the DOTS server.

1) To retrieve all the black-lists rules programmed by the DOTS client.

```
GET {scheme}://{host}:{port}/.well-known/{URI suffix for filtering}
```

2) To retrieve specific black-list rules programmed by the DOTS-cient.

```
GET {scheme}://{host}:{port}/.well-known/{URI suffix for filtering}
Accept: application/json
Content-type: application/json
 {
   "policy-id": number
 }
```

Figure 9: GET to retrieve the rules

5. IANA Considerations

TODO

6. Security Considerations

TODO

HTTPS MUST be used for data confidentiality and (D)TLS based on client certificate MUST be used for mutual authentication. The interaction between the DOTS agents requires Datagram Transport Layer Security (DTLS) and Transport Layer Security (TLS) with a ciphersuite offering confidentiality protection and the guidance given in [RFC7525] must be followed to avoid attacks on (D)TLS.

Special care should be taken in order to ensure that the activation of the proposed mechanism won't have an impact on the stability of the network (including connectivity and services delivered over that network).

Involved functional elements in the cooperation system must establish exchange instructions and notification over a secure and authenticated channel. Adequate filters can be enforced to avoid that nodes outside a trusted domain can inject request such as deleting filtering rules. Nevertheless, attacks can be initiated from within the trusted domain if an entity has been corrupted. Adequate means to monitor trusted nodes should also be enabled.

7. Acknowledgements

Thanks to C. Jacquenet for the discussion and comments.

8. References

8.1. 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>.

[RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <http://www.rfc-editor.org/info/rfc7525>.

8.2. Informative References

[REPORT] "Worldwide Infrastructure Security Report", 2014, <http://pages.arbornetworks.com/rs/arbor/images/WISR2014.pdf>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <http://www.rfc-editor.org/info/rfc791>.

[RFC4732] Handley, M., Ed., Rescorla, E., Ed., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, DOI 10.17487/RFC4732, December 2006, <http://www.rfc-editor.org/info/rfc4732>.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <http://www.rfc-editor.org/info/rfc5575>.

[RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <http://www.rfc-editor.org/info/rfc6269>.

[RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <http://www.rfc-editor.org/info/rfc7159>.

Appendix A. BGP

BGP defines a mechanism as described in [RFC5575] that can be used to automate inter-domain coordination of traffic filtering, such as what is required in order to mitigate DDoS attacks. However, support for BGP in an access network does not guarantee that traffic filtering will always be honored. Since a DOTS client will not receive an acknowledgment for the filtering request, the DOTS client should monitor and apply similar rules in its own network in cases where the DOTS server is unable to enforce the filtering rules. In addition, enforcement of filtering rules of BGP on Internet routers are usually governed by the maximum number of data elements the routers can hold as well as the number of events they are able to process in a given unit of time.

Authors' Addresses

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tireddy@cisco.com

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com

Prashanth Patil
Cisco Systems, Inc.

Email: praspati@cisco.com

Mike Geller
Cisco Systems, Inc.
3250
Florida 33309
USA

Email: mgeller@cisco.com

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Robert Moskowitz
HTT Consulting
Oak Park, MI 42837
United States

Email: rgm@htt-consult.com