

DOTS
Internet-Draft
Intended status: Informational
Expires: October 22, 2015

D. Migault, Ed.
Ericsson
April 20, 2015

DDoS Open Threat Signaling use cases
draft-mglt-dots-use-cases-00

Abstract

The document details use cases to mitigate DDoS attacks. These use cases are expected to illustrate involved communications to detect and mitigate DDoS attacks. It is expected that these communications will be in the future handled by the DDoS Open Threat Signaling (DOTS). These scenarios are intended to be useful to derive requirements for the design of DDoS Open threat Signaling.

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 October 22, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|---------------------------------------|----|
| 1. Introduction | 2 |
| 2. Terminology and Acronyms | 3 |
| 3. On-premise use case | 4 |
| 3.1. Symmetric | 5 |
| 3.2. Asymmetric | 8 |
| 4. Cloud Use Case | 10 |
| 5. Hybrid Cloud Use Case | 11 |
| 6. Security Considerations | 12 |
| 7. Privacy Considerations | 12 |
| 8. IANA Considerations | 12 |
| 9. Acknowledgments | 12 |
| 10. Normative References | 12 |
| Author's Address | 12 |

1. Introduction

DDoS is a major threat that affects any organizations of any size. In addition, these attacks have become more and more frequent, complex and sophisticated which makes DDoS attacks harder to be detected at a single point.

More specifically, traditional SYN TCP or ICMP flood attacks were relatively easy to detect at the border of the network by an on-premise device. Although such DDoS attacks remain, DDoS attacks become more and more applications specific. This results in more specialized DDoS attacks, that require a fine grained monitoring to detect suspicious traffic.

For example, DNS can be used as a channel to establish a communication channel between a bot and its Command and Control (CC) channel. A generic DNS flow traffic monitoring is not sufficient to detect such attacks. Instead it may require monitoring FQDNs with NXDOMAIN associated to behavioral traffic analysis. DNS(SEC) or NTP are used to perform DDoS reflection attacks. Detection of these attacks may involve monitoring how the source IP address may be unusually associated to heavy traffic. That said, more specific traffic monitoring and analysis is not sufficient when DDoS attacks target a specific application. In the case of slowloris flows DDoS attacks for example, the attacker initiates regular conversations with the servers, except that it maintains these conversations open. The use of TLS/DTLS makes on path monitoring impossible.

The complexity and the multitude of potential targets results in making DDoS detection a distributed system over a network. Flood attacks can be detected at the entrance of the network, SYN flood may be detected by firewalls associated to behavioral analysis. TLS and HTTP floods or low and slow and application based DDoS attacks are expected to be detected on the server side.

The multitude of DDoS monitoring appliance requires coordination. Coordination is necessary in order to manage the DDoS appliances as well as to collect the various information provided by each appliance and correlate these piece of information. Such correlation is expected to provide early detection, as well as more accurate alarms. Once a DDoS attack has been detected, the mitigation should proceed. Mitigation could be handled locally or outsourced.

The document details use cases to mitigate DDoS attacks. These use cases are expected to illustrates involved communications to detect and mitigate DDoS attacks. It is expected that these communications will be in the future handled by the DDoS Open Threat Signaling (DOTS). These scenarios are intended to be useful to derive requirements for the design of DDoS Open threat Signaling.

The document illustrates how DOTS makes possible DDoS to go beyond the scope of an isolated appliance and :

- A) Make possible a global and cross layered DDoS Monitoring, to make DDoS detection more accurate and earlier.
- B) Make possible a global and cross layer DDoS Mitigation, to mitigate in an coherent and efficient way.
- C) Make possible to share monitored information between multiple parties.
- D) Make possible to share and delegate DDoS monitoring and mitigation to third party.

2. Terminology and Acronyms

- Deny of Service (DoS): is an attack that makes resource of a service unavailable for its intended users. The resource may be computing or networking resource.
- Distributed Deny of Service (DDoS): is a DoS attack where the resources used by the attacker to perform the attack are distributed.

- DDoS Monitoring: designates the ability to inspect and monitor the traffic. This may include, exporting flow information to a Flow Repository or generating an alarm to the DDoS Controller when some threshold have been reached. In this document, DDoS Monitoring represents indifferently either a specific and dedicated DDoS Appliance, a virtual DDoS Appliance or a module.
- DDoS Mitigation: designates the ability to mitigate the DDoS attack. This may include providing filtering rules for example. In this document, DDoS Mitigation represents indifferently either a specific and dedicated DDoS Appliance, a virtual DDoS Appliance or a module.
- DDoS Controller: designates the entity that centralized monitoring, the alarms received and provides the mitigation actions. As DDoS attacks become more and more complex, a single DDoS monitoring device become dedicated to limited aspect of DDoS. As a result, these devices have only a fractional view of the ongoing activity. On the other hand, the DDoS Controller can aggregate and correlate this information have as such has a global view of the attacks. As result the DDoS Controller is more likely to take the appropriated decision to mitigate the attack.
- DDoS Appliance: designates an appliance that embeds DDoS Monitoring and/or DDoS Mitigation function. In this document, DDoS Appliance can be indifferently a hardware or virtual virtual DDoS Appliance.
- Flow Repository: designates the entity that centralized all the flow information. The Repository, may be shared between various entities and third parties. In fact, it is expected that information could be shared between independent actors, in order to mitigate DDoS Internet wild.
- Service: designates the destination of the traffic and the service that is under attack.

3. On-premise use case

The on-premise uses cases describe scenarios where DDoS is detected and mitigated on site. Section 3.1 describes the symmetric on-premise scenario, where the DDoS Appliance is place on path both the inbound and outbound traffic to the Service. Section 3.2, on the other hand presents the case where only a sub traffic is dynamically directed to the DDoS Appliance.

3.1. Symmetric

As depicted in Figure 1 the DDoS Appliance is on path of the inbound and outbound traffic to the Service. In other words, traffic coming from the Service to the end users goes also through the DDoS Appliance.

Such scenario may be associated to Small Office Home Office (SOHO) networks. In this case, the network, most likely, has a single DDoS Appliance. On the other hand, this scenario may also apply to large data center where, for example, each VM could be associated to a virtual DDoS Appliance.

The typical use case includes the following steps:

1. The DDoS Controller requests the DDoS Monitoring and DDoS Mitigation capabilities of the DDoS Appliance. Such request provides flexibility for both the DDoS Controller and the DDoS Appliances. First the DDoS Controller does not need to be tied to the DDoS Appliance, and so a single DDoS Controller may be used for various heterogeneous DDoS Appliances. Heterogeneity can be in term of vendors and/or in term of proposed capabilities. Similarly, this provides flexibility for the DDoS Appliances, as a DDoS Appliance may implement a subset of capabilities. In our example, the DDoS Controller, discovers both the DDoS Monitoring and DDoS Mitigating capabilities. DDoS Monitoring capabilities are necessary for monitoring the traffic and latter setting the alarms (see 2.). DDoS Mitigation capabilities are not mandatory to be requested here, as they are only expected to be used when the network is under attack. The reason the DDoS Controller requests those at this stage is to be able to plan its strategy for DDoS mitigation in advance instead of doing so while being under attack.
2. The DDoS Controller, then configures the appropriated capabilities on the DDoS Appliance. The configuration can typically be setting the thresholds upon which an alarm is raised by the DDoS Appliance to the DDoS Controller. Another type of setting may also be related to monitoring. DDoS Appliance may be configured to provide flow or resource (like CPU usage) information. These information may be exported to the Flow Repository in an appropriated format that enabled processing and correlation analysis by the DDoS Controller.
3. The DDoS Appliance sends the monitoring information to the Flow Repository. Note that the Flow Repository must be provided some means to authenticated the received packets as well as to

check the received information corresponds to the one requested by the DDoS Controller.

4. The DDoS Appliance raises an alarm that some suspicious traffic has been detected. This alarm corresponds to the settings performed by the DDoS Controller in step 2. As mentioned in Section 1 it may be difficult for the DDoS Appliance to determine from a local observation that a DDoS attack is ongoing or not. This is the reason the alarm is raised for suspicious traffic.
5. The DDoS Controller analyzes and correlates the received alarm for suspicious traffic and confirm or not that a DDoS attack is ongoing. Confirmation may require the DDoS Controller to perform some traffic analysis and correlates the alarm with some additional data. To do so, the DDoS Controller may consult the Flow Repository.
6. The DDoS Controller concludes that the network is under attack, and so proceeds to DDoS mitigation. In this example, the DDoS Controller is aware of the DDoS Mitigation capabilities of the DDoS Appliance as it has proceeded to the discovery mechanism in step 1. If that is not the case, the DDoS Controller should discover the DDoS mitigation capacities now. DDoS mitigations performed by the DDoS Controller are related to DDoS service. This may include for example setting some filtering rules or activation rate limitation. If traffic redirection should be performed, it is not expected to be performed by the DDoS Controller. In fact redirection implies a network reconfiguration and is considered outside the scope of the DDoS Controller. In addition to mitigate the DDoS attack, the DDoS Controller may also adjust its DDoS Monitoring settings. Motivations for doing so, may be for example to reduce the traffic on the network, or reversely, to provide a more accurate monitoring.
- 6bis. Eventually, the DDoS Controller may conclude that the network is not under attack. In this case the alarm is ignored or acknowledged to avoid the alert is re-sent and eventually load the network or the DDoS Controller. Similarly to step 6, the DDoS may also decide to adjust the monitoring settings to reduce false positive alarms. Note that the latest should be used cautiously as, such mechanism may be used as a vector of attack.

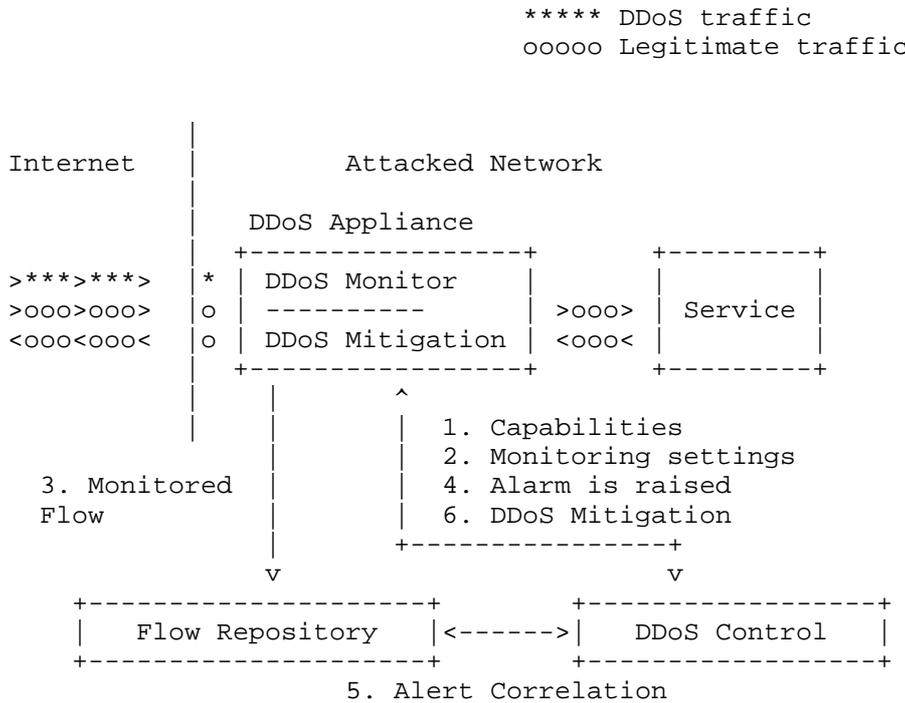


Figure 1: On-premise Symmetric Use Case

Figure 1 shows the DDoS Controller as distinct from the DDoS Appliance. In fact nothing prevents the DDoS Controller to be located on the DDoS Appliance. In this case the communications between the DDoS Controller and the DDoS Monitoring or DDoS Mitigation functions would be implementation dependent and thus outside of the scope of DOTS. The DDoS Appliance may embed a basic and limited DDoS Controller for basic configuration of the device. This is one reason why a DDoS Appliance may be configured by multiple DDoS Controllers.

Similarly, there is no requirements that the DDoS Controller belongs to the same network as the DDoS Appliances. The DDoS Controller could be placed inside the on-premise DDoS Appliances' network or remotely see Section 5 for more details.

How the DDoS Controller handles alarms and determines a suspicious traffic corresponds or not to a DDoS attack is out of scope of DOTS. Similarly, the mitigating strategies are also out of scope of DOTS.

3.2. Asymmetric

The asymmetric on-premise scenario optimize resources compared to the symmetric on-premise scenario. More specifically, in the symmetric on premise scenario, the traffic going from the Service to the end users also goes through the DDoS Appliance. Such deployment may lead to unnecessary load on the DDoS Appliance. In fact, the outbound traffic may not need to be either monitored or mitigated, and as such may reduce the packet rate or bit rate upper bound limit for inbound traffic. This may be one motivation for splitting the DDoS Monitoring module and the DDoS Mitigation modules in two different DDoS Appliances. In addition, for large networks, having a dedicated DDoS Appliance for DDoS mitigation may rationalize the cost and use of DDoS Mitigation Appliances. In fact, DDoS Mitigation Appliances may be shared by multiple Services or instances of VM of a given Service. As a result, the DDoS Mitigation Appliance do not need to scale the service traffic but instead the traffic of DDoS attacks -- which is most likely expected to remain smaller. This may not be the case for the DDoS Monitor Appliance as there is a need to always monitor the whole service traffic.

In the use case depicted by Figure 2 and Figure 3 the DDoS Mitigation Appliance only handles DDoS traffic.

The typical use case includes the following steps:

1. corresponds to the capabilities discovery phase. It is similar as the one exposed in Section 3.1. The main difference remains that DDoS Monitoring capabilities and DDoS Mitigating capabilities are discovered on two distinct DDoS Appliances.
- 2., 3., 4. and 5. corresponds to the monitoring and alarms settings. Monitoring may result in exporting data to the Flow Repository. This is similar as the steps described in Section 3.1.
6. If the DDoS Controller determines the network is under a DDoS attack, mitigation is performed in two steps. They may be ordered differently depending on criteria that are beyond the scope of this use case. First, the DDoS Mitigation is configured as described in Section 3.1 as a result of an analysis performed by the DDoS Controller. Then, traffic redirection is performed. In our case, the redirected traffic corresponds only to the inbound traffic from the end users. The traffic from the service to the end users is not redirected. This operation is not directly handled by the DDoS Controller. It can be performed manually, or upon a request from the DDoS Controller. This request is then treated by a

network management function in order to perform the appropriated network configurations.

6bis. In the case, the DDoS Controller determines the network is not under a DDoS attack, this step similar to the one described in Section 3.1.

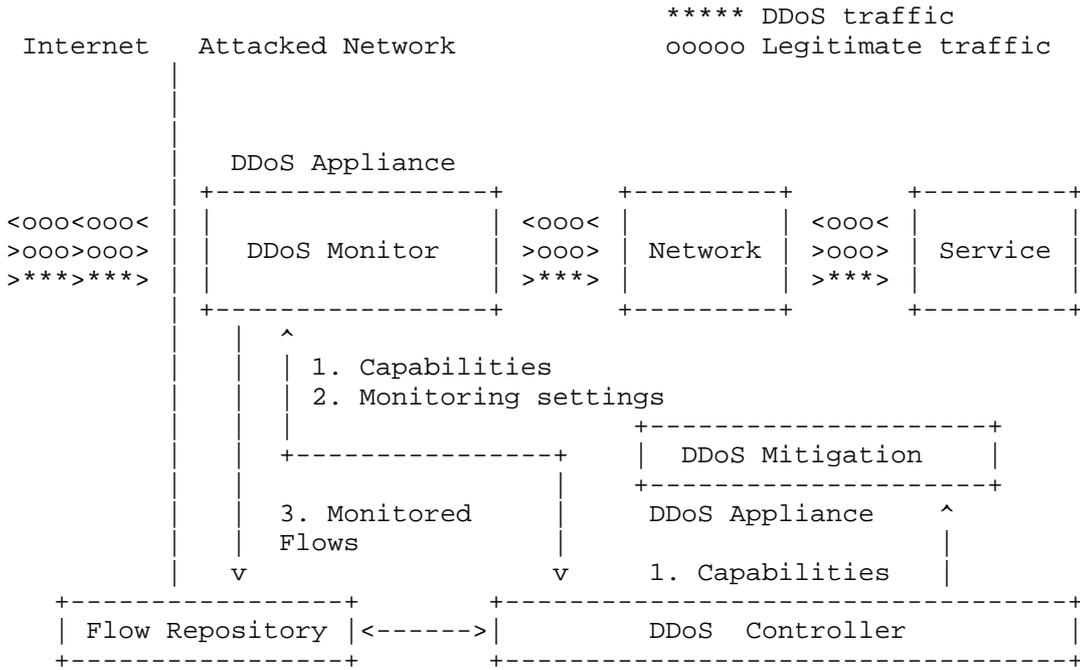


Figure 2: On-premise Asymmetric Use Case Monitoring Phase

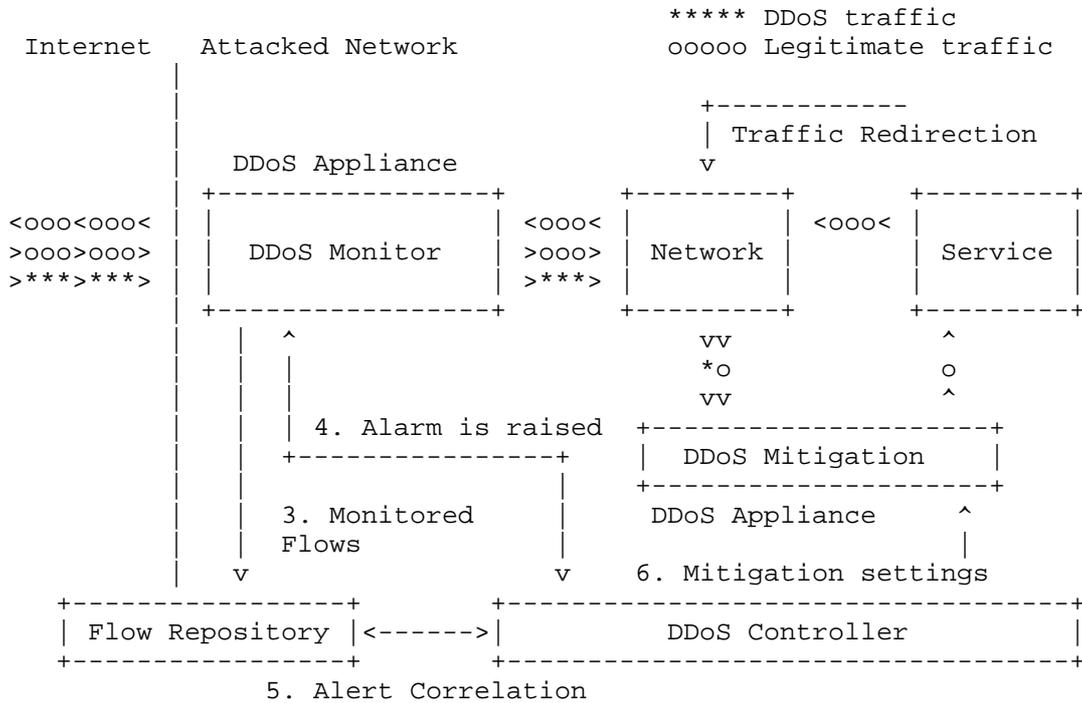


Figure 3: On-premise Asymmetric Use Case Mitigation Phase

4. Cloud Use Case

Figure 4 illustrates the Cloud use case. In this scenario, the entire DDoS monitoring and mitigation service is outsourced to a third party designated as Cloud Based DDoS Cleaning Service or Cloud for short. In order to do so, the traffic associated to the Service goes through the Cloud Based DDoS Cleaning Service as detailed in Figure 4. On the other hand, this scenario makes DDoS mitigation transparent to the Service provider, which then benefits from a "clean pipe".

Figure 4 presents the case where the Cloud is on path of both inbound and outbound traffic, a similar scenario may also consider that only the inbound traffic, that is the traffic destined to the service is directed to the cloud whereas the outbound traffic destined to the users does not.

Internal organization of the Cloud Based DDoS Cleaning Service is transparent to the Service provider. A combination of the on-premises scenarios may be used.

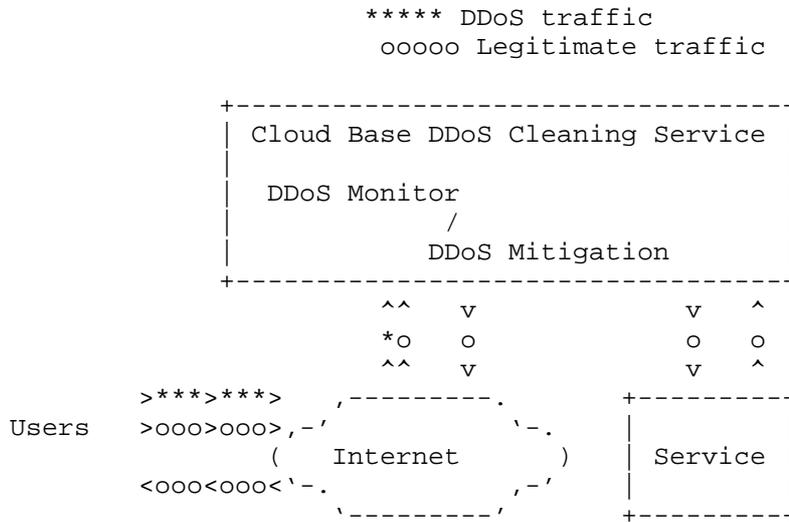


Figure 4: Cloud Use Case

5. Hybrid Cloud Use Case

The inconvenient the cloud use case scenario described in Section 4 is that redirecting the traffic to the cloud is likely to introduce additional latency. This is inconvenient as it adds a constant service degradation and cost to the Service provider. In order to address this, this section details the Hybrid Cloud scenario that combines the on-premise scenarios detailed in Section 3 and the cloud scenario detailed in Section 4

The main driver for combining the cloud and on-premise scenarios is to be able to outsource the DDoS attack mitigation to a third party only when the Service provider is under attack, or when it is not able to handle the ongoing DDoS attack. In the general case, the determination on how the service provider is able to cope and detect a DDoS attack is up to the Service provider. A continuum of scenarios can be considered and this section details only a few of them.

A specific case may consider that DDoS mitigation is outsourced by outsourcing the DDoS Controller to a third party. This DDoS Controller, drives the DDoS Monitor functions on the premise. When an alert is raised, the DDoS Controller may take the decision to mitigate internally with the DDoS attack only using on-premise facilities. This case correspond to the scenarios detailed in Section 3, except that the DDoS Controller is either located remotely, or at least accessed remotely by the third party. On the

other hand, the DDoS Controller may also decide that the DDoS attack cannot be mitigated on premise, and that mitigation should be outsourced to a cloud service as described in Section 4. In this case, the DDoS Controller is expected to redirect at least the inbound traffic of the Service provider to the cloud infrastructure. This case corresponds to the on premise asymmetric scenario detailed in Section 3.2. The difference is that redirection does not occur inside the Service provider, but involves sites redirection -- most likely using BGP signaling.

Another scenario may provide more independence to the Service provider. In this scenario, the Service provider, may have the complete control on the DDoS Monitor and DDoS Mitigation Appliances, and only uses the Cloud as a backup solution when it is not likely to deal with the DDoS attack. In this case, the DDoS Controller sends an alert to the DDoS Controller of the third party. The third party first analyzes the attack, which may require to grant access to the third party to the Flow Repository. If DDoS mitigation action are performed by the third party DDoS Controller, means should be provided to transmit information from the third party DDoS Controller to the DDoS Appliances. This could be done for example by providing access to the DDoS Appliances, or by DDoS Controller that acts as a proxy for the third party DDoS Controller.

6. Security Considerations

7. Privacy Considerations

8. IANA Considerations

This document makes no request of IANA.

9. Acknowledgments

10. Normative References

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

Author's Address

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

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