INTERNET-DRAFT
Intended Status: Standards Track
Expires: September 21, 2016

Luyuan Fang
Deepak Bansal
Microsoft

March 21, 2016

Inter-Cloud DDoS Mitigation API
draft-fang-i2nsf-inter-cloud-ddos-mitigation-api-01

Abstract

This document defines an Inter-Cloud DDoS Mitigation Abstract Layer and corresponding standardized APIs to enable the exchange of real time automated information to enable DDoS mitigation across Cloud Service Providers and Network Service Providers.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Table of Contents

1. Introduction

We recently observe the following characteristics of the DDoS attacks in the Cloud era: 1) Growing in volume: for example, 450 Gbps peak speed DDoS attack in an ISP network was observed in December 2014, while over 300 Gbps DDoS attack was reported in 2013; 2) Growing in frequency; 3) Using Cloud services to launch major attacks, especially when some cloud services do not impose bandwidth and compute resource limitation; 4) Growing in sophistication: leverage vulnerable services like NTP, DNS, and BitTorrent to amplify the available bandwidth; 5) Growing attack to Inter-cloud/Inter-provider connection links, large volume attack can disrupt all cloud services traversing through the inter-connection links.

This draft is focus on Inter-Cloud/Inter-provider DDoS attack mitigation. The fast growth in volume and scale of Distributed Denial of Service (DDoS) attacks, particularly its impact on the large pipes of Inter-Cloud, Inter-Provider connections, calls for mechanisms to enable DDoS mitigation across Cloud Service Providers (CSPs) and Network Service Providers (NSPs). These mechanisms require to define an Inter-Cloud DDoS Mitigation Abstract Layer with corresponding standardized APIs to allow real time, automated information exchange among CSPs and NSPs, and achieve rapid protective response and effective Inter Cloud/Inter Provider DDoS attack mitigation. The need for such standard Inter-Cloud DDoS Mitigation APIs is strong and urgent.

This document defines the Inter-Cloud DDoS Mitigation Abstract Layer and APIs.

This document focuses on Inter-Cloud, Inter-Provider automated exchange of DDoS Mitigation information, although similar APIs could be used within each cloud for handling malicious traffic.

1.1. Terminology

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

This document uses the terminology defined in [I-D.draft-ietf-i2nsf-gap-analysis].

In addition, this document uses the following terms.

Term	Definition
BGP	Border Gateway Protocol
CSP	Cloud Service Provider
DC	Data Center
DCI	Data Center Interconnect
DDoS	Distributed Denial of Service
DLC	Disruption Life Cycle
Inter-Cloud	The interconnection between the cloud of different providers
NSP	Network Service Provider
SDN	Software Defined Network
SVR	Server

2. Problem Statement

Along with the rapid growth of cloud services, the large pipes of Inter-Cloud, Inter-Provider connections are increasingly the subject of DDoS attacks. Since these connections are between clouds of different providers, implementing mechanism to achieve rapid protective response in case of attack is challenging. While within its own cloud each provider may be able to protect effectively its network using various DDoS protection techniques, for the Inter-Cloud/Inter-Provider links, each provider does not have full visibility of the attack, and therefore response times may be longer, counter-measures may be less effective, and therefore the severity and impact of the attacks may be very significant.

Large DDoS attacks targeting the Inter-Cloud, Inter-Provider links may consume the available bandwidth or the router/switch/server resources within tens of seconds. While the attack is on, legitimate traffic is prevented from being forwarded over the saturated links. With saturated Inter-Cloud, Inter-Provider links, even if within each cloud the DDoS mitigation may be working effectively, it can quickly be rendered irrelevant.

How does Distributed DoS attack relate to Inter-Cloud connections? The DDoS attack can be targeting the hosts, servers, end-points, gateways, or any devices in between. Regardless of the target, the attack traffic flows through the "Pipes"/inter-connection links, and can saturate these large pipes. Attack volume is the key issue here. DDoS attack BW is increasing very fast is the recent years. Attack BW greater than 100G is not uncommon any more, and 450G peek speed DDoS attack has been seen in some SP networks end of 2014. The DDoS attack can consume BW, impact multi-region Data Centers and Inter-Cloud connectivity, and interrupt multi-services. Because of its massive scale, it can also make fast mitigation more challenging.

Today, exchange of DDoS attack information and mitigation strategy among providers is largely manual and typically relies on customized operation processes established ad hoc between each provider. Manual means someone has to send emails, or make phone calls to reach the people in another Cloud, another ISP, etc. No signaling, no common API, no automation across the provider boundaries available. Because of largely manual escalation procedures, providers' reaction times to DDoS attacks to Inter-Cloud, Inter-Provider links tends to be slow (it can easily take tens of minutes if not hours to put effective mitigation measures in place) compared to Intra-Cloud DDoS mitigation, and thus the damage caused by such attacks can be substantial. The reaction time may exceed the Disruption Life Cycle (DLC) of the attack.

Sophisticated and determined malicious attackers are able to quickly learn the intended Inter-Cloud Inter-Provider link capabilities and limitations through probing. This includes bandwidth capacity, saturation resistance - the attack cannot saturate the connection links and make them unusable, and DDoS absorption resilience of the link - the attack can be absorbed without taking down the network connections and impact the services. The attacker is also able to learn the DDoS countermeasures and their response times, from which the attacker can infer the DLC that can be exacted toward the intended target. The DLC is measured by the assailant from the time the attack is initiated to the time the mitigation response becomes evident. An attacker can then use this information to design the attacks in such a way that the current and subsequent attacks inflict the most harm.

In order to achieve rapid protective response, the exchange of DDoS mitigation information between providers must be enabled in real time and in an automated, standardized fashion.

3. Inter-Cloud DDoS Mitigation Layer

The Inter-Cloud DDoS Mitigation Layer and its corresponding standardized, secure Inter-Cloud DDoS Mitigation APIs is illustrated in Figure 1.

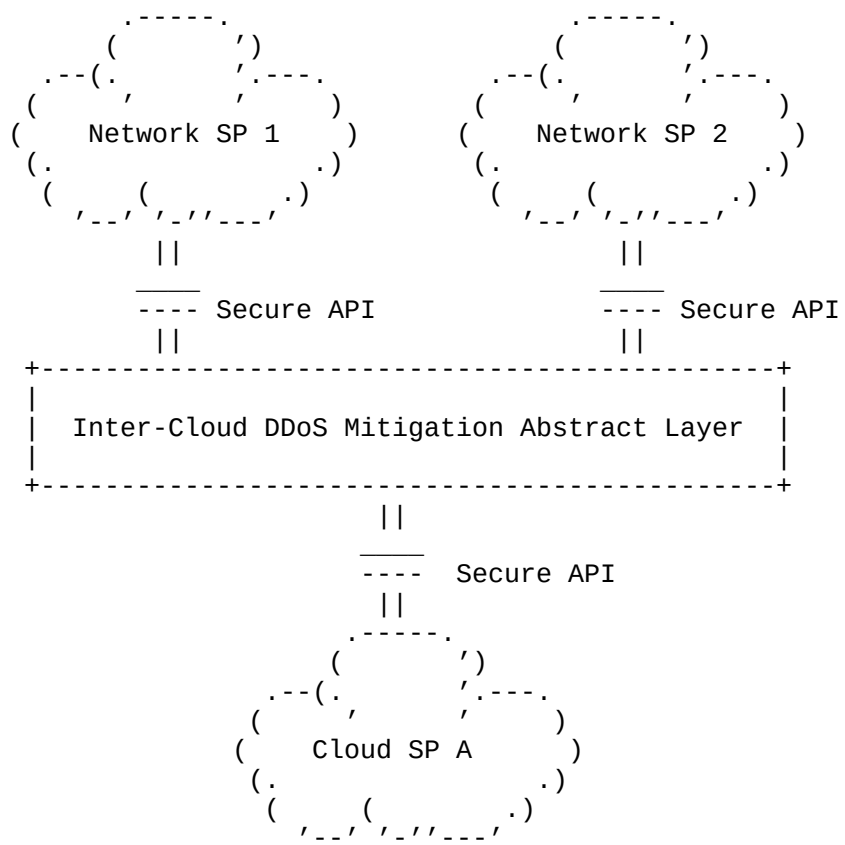

Figure 1. Inter-Cloud DDoS Mitigation Abstract Layer and APIs

Today there is no accepted industry common DDoS Mitigation Layer that can be used to reduce the reaction time and increase the effectiveness of mitigation in case of attack.

The Inter-Cloud DDoS Mitigation Abstract Layer provides standardized secure APIs that can be used by each provider to programmatically initiate real time information exchanges to other providers to provide visibility of the attack and coordinate DDoS mitigation mechanisms, Exchanged information may include signatures and forensic of the attack, timestamps, and black-holing countermeasures.

The Inter-Cloud DDoS Mitigation Abstract Layer provides corresponding API calls to exchange mitigation information on the following areas.

DDoS Protection Types:

- TCP flood rate limiting
- UDP flood rate limiting
- TCP SYN.ACK/RST flood protection and authentication
- Maximum concurrent connections per interval rate limiting
- Maximum number of new connections allowed per interval rate limiting
- Maximum fragment packets allowed per interval rate limiting
- Maximum number of packets allowed per interval rate limiting
- Black-holing
- Use BGP Flowspec [RFC5575] to auto-coordinate traffic filtering, DDoS mitigation
- Other BGP Signaling and Mitigation examples
  - BGP /24 route advertisement with community string option
  - Mitigation support for /32 with type and rate limit thresholds
  - /32 removal from mitigation
  - BGP support for /24 removal

Attack Lifecycle Monitoring and Reporting

- Volume and scale of the attack, signatures, forensic
- Timestamps

4. Inter-Cloud DDoS Mitigation API

4.1. Categories of Inter-cloud API

The following describe the basic functions the Inter-Cloud DDoS mitigation MUST support.

4.1.1. Capability information exchange:

Support "Query" the DDoS capabilities from one provider to another provider.

4.1.2. Mitigation Request and response:

Mitigation Request: One provider can "Request" for mitigation by partner provider based on pre-agreement.

Mitigation Response: The provider received DDoS mitigation request first acknowledge the request, then execute a particular DDoS capability on behalf of the requesting provider, and respond back with the logged actions performed and mitigation status.

4.1.3. Monitoring and Reporting:

Monitoring: Allow another provider to monitor DDoS status and mitigation processes.

Reporting: Provider DDoS status reports to partner providers.

4.1.4. Knowledge sharing:

Allow partner providers to query for a specific DDoS related data to enhance their DDoS resiliency and perform coordinate mitigation when possible.

4.2. REST API format

4.2.1. Capability

Definition: A participating provider should allow another provider to query for its DDoS capabilities.

The following REST API are the basic ones that every provider participating MUST provide.

4.2.1.1. GET

Example 1: GET (DDoS mitigation Capabilities)

a. Description: The receiving provide returns a list of DDoS mitigation it can perform

b. Parameters: None

c. Responses: 200, an array of mitigation objects format.

Example 2: GET (DDoS mitigation Capabilities - protocol)

a. Description: Return a list of DDoS mitigation that this provider can perform for the protocol specified.

b. Parameters: protocol is one of the following strings {tcp, udp, dns}

c. Responses: 200, OK, an array of mitigation objects format.

(more details to be added especially around format of the object to be returned).

4.2.2. Mitigation

Definition: Mitigation Request and Response must be supported between participating providers for executing a particular DDoS capability.

The following REST API are the baselines that each participating providers MUST support.

4.2.2.1. POST

a. Description: Create a new policy what will cause a mitigation to be performed based on a specific trigger.

b. Parameters: PolicyObject {To be specified}

c. Responses: 200, OK, return an identifier.

4.2.2.2. GET

a. Description: Get an existing policy.

b. Parameters: id identifier of the policy that was created.

c. Responses: 200, OK, return the policy of the specified id

4.2.2.3. PUT

a. Description: Update a particular policy.

b. Parameters: PolicyObject {To be specified} & id which is the identifier which was returned after a successful create of a policy.

4.2.2.4. DELETE

a. Description: Delete a policy and therefore end any mitigation that is currently active.

b. Parameters: id the identifier of the policy that was created.

c. Response: 200, OK, policy deleted.

4.2.3. Monitor & Reporting

Definition: A participating provider MUST allow another provider to monitor a particular DDoS mitigation.

The following REST API are the basic ones that every provider must provide.

4.2.3.1. POST

a. Description: Created a new monitored object for a policy/mitigation.

b. Parameter: MonitoredObject {To be specified} & id which is the mitigation identifier. The MonitoredObject will have parameter to enable retrieving sFlow to a particular endpoint for collection of the metrics. By default, you can use REST API calls as defined below to retrieve monitored objects stats.

c. Responses: 200, OK

4.2.3.2. GET

a. Description: Get the current monitoring settings for this mitigation/policy.

b. Parameter: id identifier for the mitigation/policy.

c. Responses: 200, OK, monitoring settings

4.2.3.3. PUT

a. Description: Update the current monitoring settings for this mitigation/policy

b. Parameter: id identifier for the mitigation/policy.

c. Responses: 200, OK

4.2.3.4. DELETE

a. Description: Remove all monitoring configuration for this mitigation/policy

b. Parameter: id identifier for the mitigation/policy.

c. Responses: 200, OK

4.2.3.5. GET

a. Description: Return the stats available for this mitigation and monitored object.

b. Parameters: None

c. Responses: 200, OK, stats

4.2.4. Knowledge Sharing

Definition: A participating provider MUST allow another participating provider to query for a specific DDoS related data to enhance their DDoS resiliency.

The following REST API are the basic ones that every provider must provide.

4.2.4.1. GET

a. Description: Return the current blacklist.

b. Parameter: Size to limit the returned list.

c. Responses: 200, OK, return a string array of blacklisted IPs.

5. Security Considerations

Given the subject of the draft is Inter-Cloud/Inter-Provider DDoS mitigation, security policies among the participating providers must be agreed upon and strictly followed. Authentication MUST be enforced on all interconnections and APIs in discussion.

6. IANA Considerations

None.

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.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules," RFC

5575, August 2009.

7.2. Informative References

[I-D.draft-ietf-i2nsf-gap-analysis] S. Hares et al., "Analysis of Use Cases and Gaps in Technology for I2NSF ",draft-ietf-i2nsf-gap-analysis-00.txt, Feb. 2016.

Authors' Addresses

Luyuan Fang
Microsoft
15590 NE 31st St
Redmond, WA 98052
Email: lufang@microsoft.com

Deepak Bansal
Microsoft
15590 NE 31st St
Redmond, WA 98052
Email: dbansal@microsoft.com

Contributing Authors' Addresses

Jim Nyland
Microsoft
15590 NE 31st St
Redmond, WA 98052
Email: jnyland@microsoft.com

Geoff Outhred
Microsoft
15590 NE 31st St
Redmond, WA 98052
Email: geoffo@microsoft.com

Anh Cao
Microsoft
15590 NE 31st St
Redmond, WA 98052
Email: anhcao@microsoft.com