

Behave
Internet-Draft
Intended status: Standards Track
Expires: July 18, 2013

S. Sivakumar
R. Penno
Cisco Systems
January 14, 2013

IPFIX Information Elements for logging NAT Events
draft-sivakumar-behave-nat-logging-06

Abstract

NAT devices are required to log events like creation and deletion of translations and information about the resources it is managing. With the wide deployment of Carrier Grade NAT (CGN) devices, the logging of events have become very important for legal purposes. The logs are required in many cases to identify an attacker or a host that was used to launch malicious attacks and/or for various other purposes of accounting. Since there is no standard way of logging this information, different NAT devices behave differently and hence it is difficult to expect a consistent behavior. The lack of a consistent way makes it difficult to write the collector applications that would receive this data and process it to present useful information. This document describes the information that is required to be logged by the NAT devices.

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 July 18, 2013.

Copyright Notice

Copyright (c) 2013 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. Terminology	3
2. Introduction	3
2.1. Requirements Language	3
3. Scope	3
4. Event based logging	4
4.1. Information Elements	4
4.2. Definition of NAT Events	7
4.3. Quota exceeded - Sub Event types	8
4.4. Templates for NAT Events	8
4.4.1. NAT44 create and delete session event	8
4.4.2. NAT64 create and delete session event	9
4.4.3. NAT44 BIB create and delete event	10
4.4.4. NAT64 BIB create and delete event	10
4.4.5. Addresses Exhausted event	10
4.4.6. Ports Exhausted event	11
4.4.7. Quota exceeded	11
4.4.8. Address Binding	12
4.4.9. Port block allocation and de-allocation	12
5. Encoding	12
5.1. IPFIX	13
6. Acknowledgements	13
7. IANA Considerations	13
8. Security Considerations	13
9. References	13
9.1. Normative References	13
9.2. Informative References	14
Authors' Addresses	14

1. Terminology

The usage of the term "NAT device" in this document refer to any NAT44 and NAT64 devices. The usage of the term "collector" refers to any device that receives the binary data from a NAT device and converts that into meaningful information. This document uses the term "Session" as it is defined in [RFC2663] and the term BIB as it is defined in [RFC6146]

2. Introduction

This document details the IPFIX Information Elements(IEs) that are required for logging by a NAT device. The document will specify the format of the IE's that are required to be logged by the NAT device and all the optional fields. The fields specified in this document are gleaned from [RFC4787] and [RFC5382].

2.1. Requirements Language

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. Scope

This document provides the information model to be used for logging the NAT devices including Carrier Grade NAT (CGN) events. This document focuses exclusively on the specification of IPFIX IE's. This document does not provide guidance on the transport protocol like TCP, UDP or SCTP that is to be used to log NAT events. The log events SHOULD NOT be lost but the choice of the actual transport protocol is beyond the scope of this document.

The existing IANA IPFIX Information Elements registry [IPFIX-IANA] already has assignments for many NAT logging events. For convenience, this document uses those same Information Elements. However, as stated earlier, this document is not defining IPFIX or Netflow 9 as the framework for logging. Rather, the information contained in these elements is within the scope of this document.

This document assumes that the NAT device will use the existing IPFIX framework to send the log events to the collector. This would mean that the NAT device will specify the template that it is going to use for each of the events. The templates can be of varying length and there could be multiple templates that a NAT device could use to log the events.

The implementation details of the collector application is beyond the scope of this document.

The optimization of logging the NAT events are left to the implementation and are beyond the scope of this document.

4. Event based logging

An event in a NAT device can be viewed as a happening as it relates to the management of NAT resources. The creation and deletion of NAT sessions and bindings are examples of events as it results in the resources (addresses and ports) being allocated or freed. The events can happen either through the processing of data packets flowing through the NAT device or through an external entity installing policies on the NAT router or as a result of an asynchronous event like a timer. The list of events are provided in Section 4.1. Each of these events SHOULD be logged, unless they are administratively prohibited. A NAT device MAY log these events to multiple collectors if redundancy is required. The network administrator will specify the collectors to which the log records are to be sent.

A collector may receive NAT events from multiple CGN devices and should be able to distinguish between the devices. Each CGN device should have a unique source ID to identify themselves. The source ID is part of the IPFIX template and data exchange.

Prior to logging any events, the NAT device MUST send the template of the record to the collector to advertise the format of the data record that it is using to send the events. The templates can be exchanged as frequently as required given the reliability of the connection. There SHOULD be a configurable timer for controlling the template refresh. NAT device SHOULD combine as many events as possible in a single packet to effectively utilize the network bandwidth.

4.1. Information Elements

The templates could contain a subset of the Information Elements (IEs) shown in Table 1 depending upon the event being logged. For example a NAT44 session creation template record will contain,

```
{sourceIPv4Address, postNATSourceIPv4Address, destinationIPv4Address,  
postNATDestinationIPv4Address, sourceTransportPort,  
postNAPTSourceTransportPort, destinationTransportPort,  
postNAPTDestTransportPort, natOriginatingAddressRealm, natEvent,  
timeStamp}
```

An example of the actual event data record is shown below - in a readable form

```
{192.168.16.1, 201.1.1.100, 207.85.231.104, 207.85.231.104, 14800,  
1024, 80, 80, 0, 1, 09:20:10:789}
```

A single NAT device could be exporting multiple templates and the collector should support receiving multiple templates from the same source.

The following is the table of all the IE's that a CGN device would need to export the events. The formats of the IE's and the IPFIX IDs are listed below.

Field Name	Size (bits)	IANA IPFIX ID	Description
timeStamp	64	323	System Time when the event occurred.
vlanID	16	58	VLAN ID in case of overlapping networks
ingressVRFID	32	234	VRF ID in case of overlapping networks
sourceIPv4Address	32	8	Source IPv4 Address
postNATSourceIPv4Address	32	225	Translated Source IPv4 Address
protocolIdentifier	8	4	Transport protocol
sourceTransportPort	16	7	Source Port
postNAPTsourceTransportPort	16	227	Translated Source port
destinationIPv4Address	32	12	Destination IPv4 Address
postNATDestinationIPv4Address	32	226	Translated IPv4 destination address
destinationTransportPort	16	11	Destination port
postNAPTdestinationTransportPort	16	228	Translated Destination port
sourceIPv6Address	27	128	Source IPv6 address
destinationIPv6Address	128	28	Destination IPv6 address

postNATSourceIPv6Address	128	281	Translated source IPv6 addresss
postNATDestinationIPv6Address	128	282	Translated Destination IPv6 address
natOriginatingAddressRealm	8	229	Address Realm
natEvent	8	230	Type of Event
portRangeStart	16	361	Allocated port block start
portRangeEnd	16	362	Allocated Port block end
portRangeStepSize	16	363	Step size of next port
portRangeNumPorts	16	364	Number of ports

Table 1: Template format Table

4.2. Definition of NAT Events

The following are the list of NAT events and the proposed event values. The list can be expanded in the future as necessary. The data record will have the corresponding natEvent value to identify the event that is being logged.

Event Name	Values
NAT44 Session create	1
NAT44 Session delete	2
NAT Addresses exhausted	3
NAT64 Session create	4
NAT64 Session delete	5
NAT44 BIB create	6
NAT44 BIB delete	7
NAT64 BIB create	8
NAT64 BIB delete	9
NAT ports exhausted	10
Quota exceeded	11
Address Binding	12
Port block allocation	13
Port block de-allocation	14

Table 2: NAT Event ID table

4.3. Quota exceeded - Sub Event types

The following table shows the sub event types for the Quota exceeded event

Quota Exceeded Event Name	Values
Max Session entries	1
Max BIB entries	2
Max entries per user	3

Table 3: Sub Event ID table

4.4. Templates for NAT Events

The following is the template of events that will have to be logged. The events below are identified at the time of this writing but the events are expandable. Depending on the implementation and configuration various IE's specified can be included or ignored.

4.4.1. NAT44 create and delete session event

This event will be generated when a NAT44 session is created or deleted. The template will be the same, the natEvent will indicate whether it is a create or a delete event. The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
vlanID/ingressVRFID	32	No
sourceIPv4Address	32	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	Yes
sourceTransportPort	16	Yes
postNAPTsourceTransportPort	16	Yes
destinationIPv4Address	32	No
postNATDestinationIPv4Address	32	No
destinationTransportPort	16	No
postNAPTdestinationTransportPort	16	No
natOriginatingAddressRealm	8	No
natEvent	8	Yes

Table 4: NAT44 Session delete/create template

4.4.2. NAT64 create and delete session event

This event will be generated when a NAT64 session is created. The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
vlanID/ingressVRFID	32	No
sourceIPv6Address	128	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	Yes
sourceTransportPort	16	Yes
postNAPTsourceTransportPort	16	Yes
destinationIPv6Address	128	No
postNATDestinationIPv4Address	32	No
destinationTransportPort	16	No
postNAPTdestinationTransportPort	16	No
natOriginatingAddressRealm	8	No
natEvent	8	Yes

Table 5: NAT64 session create/delete event template

4.4.3. NAT44 BIB create and delete event

This event will be generated when a NAT44 Bind entry is created. The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
vlanID/ingressVRFID	32	No
sourceIPv4Address	32	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	No
sourceTransportPort	16	No
postNAPTsourceTransportPort	16	No
natOriginatingAddressRealm	8	No
natEvent	8	Yes

Table 6: NAT44 BIB create/delete event template

4.4.4. NAT64 BIB create and delete event

This event will be generated when a NAT64 Bind entry is created. The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
vlanID/ingressVRFID	32	No
sourceIPv6Address	128	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	No
sourceTransportPort	16	No
postNAPTsourceTransportPort	16	No
natOriginatingAddressRealm	8	No
natEvent	8	Yes

Table 7: NAT64 BIB create/delete event template

4.4.5. Addresses Exhausted event

This event will be generated when a NAT device runs out of global IPv4 addresses in a given pool of addresses. Typically, this event would mean that the NAT device wont be able to create any new translations until some addresses/ports are freed. The following is

a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natEvent	8	Yes
natPoolName	String	Yes

Table 8: NAT Address Exhausted event template

4.4.6. Ports Exhausted event

This event will be generated when a NAT device runs out of ports for a global IPv4 address. Port exhaustion shall be reported per protocol (UDP, TCP etc) The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natEvent	8	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	Yes

Table 9: NAT Ports Exhausted event template

4.4.7. Quota exceeded

This event will be generated when a NAT device cannot allocate resources as a result of an administratively defined policy. The examples of Quota exceeded are to allow only certain number of NAT sessions per device, certain number of NAT sessions per user etc. The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natEvent	8	Yes
natLimitEvent	32	Yes
sourceIPv4 address	32	No
sourceIPv6 address	128	No

Table 10: NAT Quota Exceeded event template

4.4.8. Address Binding

This event will be generated when a NAT device binds a local address with a global address. This binding event happens when the first packet of the first flow from a host in the private realm.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natEvent	8	Yes
sourceIPv4 address	32	No
sourceIPv6 address	128	No
Translated Source IPv4 Address	32	8

Table 11: NAT Address Binding template

4.4.9. Port block allocation and de-allocation

This event will be generated when a NAT device allocates/de-allocates ports in a bulk fashion, as opposed to allocating a port on a per flow basis. NAT devices would do this in order to reduce logs and potentially to limit the number of connections a subscriber is allowed to use. In the following Port Block allocation template, the portRangeStart must be specified. Along with portRangeStart, at least one of portRangeEnd, portRangeStepSize or portRangeNumPorts MUST be specified. If portRangeEnd is specified, it MUST NOT be lesser than portRangeStart. The value of portRangeStepSize MUST be between 1 and 32K.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
portRangeStart	16	Yes
portRangeEnd	16	No
portRangeStepSize	16	No
portRangeNumPorts	16	No

Table 12: NAT Port Block Allocation event template

5. Encoding

5.1. IPFIX

This document uses IPFIX as the encoding mechanism to describe the logging of NAT events. However, the information that should be logged SHOULD be the same irrespective of what kind of encoding scheme is used. IPFIX is chosen because is it an IETF standard that meets all the needs for a reliable logging mechanism. IPFIX provides the flexibility to the logging device to define the data sets that it is logging. The information elements specified for logging MUST be the same irrespective of the encoding mechanism used.

6. Acknowledgements

Thanks to Dan Wing, Selvi Shanmugam, Mohamed Boucadir, Jacni Qin Ramji Vaithianathan, Simon Perreault, Jean-Francois Tremblay and Julia Renouard for their review and comments.

7. IANA Considerations

8. Security Considerations

None.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.
- [RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, October 2008.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6

Clients to IPv4 Servers", RFC 6146, April 2011.

9.2. Informative References

- [NAT-EVENT-LOG-IANA]
IANA, "NAT event log entities", 2012, <<http://www.iana.org/assignments/nat-event-log/nat-event-log.xml>>.
- [RFC5101] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.
- [RFC5102] Quittek, J., Bryant, S., Claise, B., Aitken, P., and J. Meyer, "Information Model for IP Flow Information Export", RFC 5102, January 2008.
- [RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, March 2009.

Authors' Addresses

Senthil Sivakumar
Cisco Systems
7100-8 Kit Creek Road
Research Triangle Park, North Carolina 27709
USA

Phone: +1 919 392 5158
Email: ssenthil@cisco.com

Renaldo Penno
Cisco Systems
170 W Tasman Drive
San Jose, California 95035
USA

Phone:
Email: repenno@cisco.com

