Behave
Internet-Draft
Intended status: Standards Track
Expires: December 2, 2016

S. Sivakumar
R. Penno
Cisco Systems

May 31, 2016

IPFIX Information Elements for logging NAT Events
draft-ietf-behave-ipfix-nat-logging-09

Abstract

Network operators require NAT devices to log events like creation and deletion of translations and information about the resources that the NAT device is managing. The logs are essential in many cases to identify an attacker or a host that was used to launch malicious attacks and for various other purposes of accounting. Since there is no standard way of logging this information, different NAT devices log the information using proprietary formats and hence it is difficult to expect a consistent behavior. The lack of a consistent way to log the data makes it difficult to write the collector applications that would receive this data and process it to present useful information. This document describes the formats for logging of NAT events.

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 December 2, 2016.

Table of Contents

1. Terminology 3
2. Introduction 3
2.1. Requirements Language 4
3. Scope 4
4. Deployment 4
5. Event based logging 5
5.1. Logging of destination information 6
5.2. Information Elements 6
5.3. Definition of NAT Events 8
5.4. Quota exceeded Event types 9
5.5. Threshold reached Event types 10
5.6. Templates for NAT Events 10
5.6.1. NAT44 create and delete session events 11
5.6.2. NAT64 create and delete session events 11
5.6.3. NAT44 BIB create and delete events 12
5.6.4. NAT64 BIB create and delete events 13
5.6.5. Addresses Exhausted event 13
5.6.6. Ports Exhausted event 14
5.6.7. Quota exceeded events 14
5.6.7.1. Maximum session entries exceeded 14
5.6.7.2. Maximum BIB entries exceeded 15
5.6.7.3. Maximum entries per user exceeded 15
5.6.7.4. Maximum active host or subscribers exceeded . . . 15
5.6.7.5. Maximum fragments pending reassembly exceeded . . 16
5.6.8. Threshold reached events 16
5.6.8.1. Address pool high or low threshold reached . . . 16
5.6.8.2. Address and port high threshold reached 17
5.6.8.3. Per-user Address and port high threshold reached 17
5.6.8.4. Global Address mapping high threshold reached . . 18
5.6.9. Address binding create and delete events 18
5.6.10. Port block allocation and de-allocation 19
6. Encoding 20
6.1. IPFIX 20
7. Acknowledgements 20
8. IANA Considerations 20
8.1. Information Elements 20
8.1.1. natInstanceID 20
8.1.2. internalAddressRealm 21

8.1.3. externalAddressRealm 21
8.1.4. natQuotaExceededEvent 21
8.1.5. natThresholdEvent 22
8.1.6. natEvent 23
9. Management Considerations 24
9.1. Ability to collect events from multiple NAT devices . . . 25
9.2. Ability to suppress events 25
10. Security Considerations 25
11. References 25
11.1. Normative References 25
11.2. Informative References 26
Authors' Addresses 27

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 Binding Information Base (BIB) as it is defined in [RFC6146]. The usage of the term Information Element (IE) is defined in [RFC7011]. The term Carrier Grade NAT refers to a large scale NAT device as described in [RFC6888]

The IPFIX Information Elements that are NAT specific are created with NAT terminology. In order to avoid creating duplicate IEs, IEs are reused if they convey the same meaning. This document uses the term timestamp for the Information element which defines the time when an event is logged, this is the same as IPFIX term observationTimeMilliseconds as described in [IPFIX-IANA]. Since observationTimeMilliseconds is not self explanatory for NAT implementors, this document uses the term timeStamp.

2. Introduction

The IPFIX Protocol [RFC7011] defines a generic push mechanism for exporting information and events. The IPFIX Information Model [IPFIX-IANA] defines a set of standard IEs which can be carried by the IPFIX protocol. This document details the IPFIX Information Elements(IEs) that MUST be logged by a NAT device that supports NAT logging using IPFIX, and all the optional fields. The fields specified in this document are gleaned from [RFC4787] and [RFC5382].

This document and [I-D.behave-syslog-nat-logging] are written in order to standardize the events and parameters to be recorded, using IPFIX [RFC7011] and SYSLOG [RFC5424]respectively. The intent is to

provide a consistent way to log information irrespective of the mechanism that is used.

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 events including Carrier Grade NAT (CGN) events. This document focuses exclusively on the specification of IPFIX IEs. [RFC7011] provides guidance on the choices of the transport protocols used for IPFIX and their effects. 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 IEs registry [IPFIX-IANA] already has assignments for most of the NAT logging events. This document uses the allocated IPFIX IEs and will request IANA for the ones that are defined in this document but not yet allocated.

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 is left to the implementation and is beyond the scope of this document.

4. Deployment

NAT logging based on IPFIX uses binary encoding and hence is very efficient. IPFIX based logging is recommended for environments where a high volume of logging is required, for example, where per-flow logging is needed or in case of Carrier Grade NAT. However, IPFIX based logging requires a collector that processes the binary data and requires a network management application that converts this binary data to a human readable format.

A collector may receive NAT events from multiple CGN devices. The collector distinguishes between the devices using the source IP address, source port, and Observation Domain ID in the IPFIX header.

A collector may have scale issues if it is overloaded by a large number of simultaneous events. An appropriate throttling mechanism shall be used to handle the oversubscription.

The logs that are exported can be used for a variety of reasons. An use case is to do accounting based on when the users logged on and off. The translation will be installed when the user logs on and removed when the user logs off. These events create log events. Another use case is to identify an attacker or a host in a provider network. The network administrators can use these logs to identify the usage patterns, need for additional IP addresses etc. The deployment of NAT logging is not limited to just these cases.

5. Event based logging

An event in a NAT device can be viewed as a state transition as it relates to the management of NAT resources. The creation and deletion of NAT sessions and bindings are examples of events as they result in resources (addresses and ports) being allocated or freed. The events can happen 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 Table 2. 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. The list of collectors and its associated information like the transport address, port and protocol MUST be preserved across reboots.

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. The IPFIX template management is described in detail in Section 8 of [RFC7011]. The NAT device SHOULD combine as many events as possible in a single packet to effectively utilize the network bandwidth.

5.1. Logging of destination information

Logging of destination information in a NAT event has been discussed in [RFC6302] and [RFC6888]. Logging of destination information increases the size of each record and increases the need for storage considerably. It increases the number of log events generated because when the same user connects to a different destination, it results in a log record per destination address. Logging of destination information also results in the loss of privacy and hence should be done with caution. However, this draft provides the necessary fields to log the destination information in cases where they should be logged.

5.2. Information Elements

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

{sourceIPv4Adress, postNATSourceIPv4Address, destinationIpv4Address, postNATDestinationIPv4Address, sourceTransportPort, postNAPTSourceTransportPort, destinationTransportPort, postNAPTDestTransportPort, internalAddressRealm, natEvent, timeStamp}

An example of the actual event data record is shown below - in a human 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 MUST support receiving multiple templates from the same source.

The following is the table of all the IEs that a NAT device would need to export the events. The formats of the IEs and the IPFIX IDs are listed below. Some of the IPFIX IEs are not assigned yet, and hence the detailed description of these fields are requested in the IANA considerations section.

Field Name	Size (bits)	IANA IPFIX ID	Description
timeStamp	64	323	System Time when the event occured.
natInstanceId	32	TBD	NAT Instance Identifier
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	128	27	Source IPv6 address

destinationIPv6Address	128	28	Destination IPv6 address	
postNATSourceIPv6Address	128	281	Translated source IPv6 addresss	
postNATDestinationIPv6Address	128	282	Translated Destination IPv6 address	
internalAddressRealm	8	TBD	Source Address Realm	
externalAddressRealm	8	TBD	Destination Address Realm	
natEvent	8	230	Type of Event	
portRangeStart	16	361	Allocated port block start	
portRangeEnd	16	362	Allocated Port block end	
natPoolID	32	283	NAT pool Identifier	
natQuotaExceededEvent	32	TBD	Limit event identifier	
natThresholdEvent	32	TBD	Threshold event identifier	

Table 1: Template format Table

5.3. 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
NAT Addresses exhausted	3
NAT44 Session create	4
NAT44 Session delete	5
NAT64 Session create	6
NAT64 Session delete	7
NAT44 BIB create	8
NAT44 BIB delete	9
NAT64 BIB create	10
NAT64 BIB delete	11
NAT ports exhausted	12
Quota exceeded	13
Address binding create	14
Address binding delete	15
Port block allocation	16
Port block de-allocation	17
Threshold reached	18

Table 2: NAT Event ID table

5.4. Quota exceeded Event types

The Quota exceeded events are generated when the hard limits set by the administrator has been reached or exceeded. The following table shows the sub event types for the Quota exceeded or limits reached event. The events that can be reported are the Maximum session entries limit reached, Maximum BIB entries limit reached, Maximum (session/BIB) entries per user limit reached, Maximum active hosts limit reached or maximum subscribers limit reached and Maximum Fragments pending reassembly limit reached.

Quota Exceeded Event Name	Values
Maximum Session entries	1
Maximum BIB entries	2
Maximum entries per user	3
Maximum active hosts or subscribers	4
Maximum fragments pending reassembly	5

Table 3: Quota Exceeded event table

5.5. Threshold reached Event types

The following table shows the sub event types for the threshold reached event. The administrator can configure the thresholds and whenever the threshold is reached or exceeded, the corresponding events are generated. The main difference between Quota Exceeded and the Threshold reached events is that, once the Quota exceeded events are hit, the packets are dropped or mappings wont be created etc, whereas, the threshold reached events will provide the operator a chance to take action before the traffic disruptions can happen. A NAT device can choose to implement one or the other or both.

The address pool high threshold event will be reported when the address pool reaches a high water mark as defined by the operator. This will serve as an indication that the operator might have to add more addresses to the pool or an indication that the subsequent users may be denied NAT translation mappings.

The address and port mapping high threshold event is generated, when the number of ports in the configured address pool has reached a configured threshold.

The per-user address and port mapping high threshold is generated when a single user uses more address and port mapping than a configured threshold.

Threshold Exceeded Event Name	Values
Address pool high threshold event	1
Address pool low threshold event	2
Address and port mapping high threshold event	3
Address and port mapping per user high threshold event	4
Global Address mapping high threshold event	5

Table 4: Threshold event table

5.6. Templates for NAT Events

The following is the template of events that will be logged. The events below are identified at the time of this writing but the set of events is extensible. Depending on the implementation and configuration various IEs specified can be included or ignored.

5.6.1. NAT44 create and delete session events

These events 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.

The destination address and port information is optional as required by [RFC6888]. However, when the destination information is suppressed, the session log event contains the same information as the BIB event. In such cases, the NAT device SHOULD NOT send both BIB and session events.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
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
internalAddressRealm	8	No
externalAddressRealm	8	No
natEvent	8	Yes

Table 5: NAT44 Session delete/create template

5.6.2. NAT64 create and delete session events

These events will be generated when a NAT64 session is created or deleted. The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
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
internalAddressRealm	8	No
externalAddressRealm	8	No
natEvent	8	Yes

Table 6: NAT64 session create/delete event template

5.6.3. NAT44 BIB create and delete events

These events will be generated when a NAT44 Bind entry is created or deleted. The following is a template of the event.

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

Table 7: NAT44 BIB create/delete event template

5.6.4. NAT64 BIB create and delete events

These events will be generated when a NAT64 Bind entry is created or deleted. The following is a template of the event.

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

Table 8: NAT64 BIB create/delete event template

5.6.5. 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 won't be able to create any new translations until some addresses/ports are freed. This event SHOULD be rate limited as many packets hitting the device at the same time will trigger a burst of addresses exhausted events.

The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natPoolID	32	Yes

Table 9: Address Exhausted event template

5.6.6. 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). This event SHOULD be rate limited as many packets hitting the device at the same time will trigger a burst of port exhausted events.

The following is a template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	Yes

Table 10: Ports Exhausted event template

5.6.7. Quota exceeded events

This event will be generated when a NAT device cannot allocate resources as a result of an administratively defined policy. The quota exceeded event templates are described below.

5.6.7.1. Maximum session entries exceeded

The maximum session entries exceeded event is generated when the administratively configured NAT session limit is reached. The following is the template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natQuotaExceededEvent	32	Yes
configuredLimit	32	Yes

Table 11: Session Entries Exceeded event template

5.6.7.2. Maximum BIB entries exceeded

The maximum BIB entries exceeded event is generated when the administratively configured BIB entry limit is reached. The following is the template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natQuotaExceededEvent	32	Yes
configuredLimit	32	Yes

Table 12: BIB Entries Exceeded event template

5.6.7.3. Maximum entries per user exceeded

This event is generated when a single user reaches the administratively configured NAT translation limit. The following is the template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natQuotaExceededEvent	32	Yes
configuredLimit	32	Yes
vlanID/ingressVRFID	32	No
sourceIPv4 address	32	Yes for NAT44
sourceIPv6 address	128	Yes for NAT64

Table 13: Per-user Entries Exceeded event template

5.6.7.4. Maximum active host or subscribers exceeded

This event is generated when the number of allowed hosts or subscribers reaches the administratively configured limit. The following is the template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natQuotaExceededEvent	32	Yes
configuredLimit	32	Yes

Table 14: Maximum hosts/subscribers Exceeded event template

5.6.7.5. Maximum fragments pending reassembly exceeded

This event is generated when the number of fragments pending reassembly reaches the administratively configured limit. The following is the template of the event.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natQuotaExceededEvent	32	Yes
configuredLimit	32	Yes
internalAddressRealm	8	Yes
vlanID/ingressVRFID	32	No
sourceIPv4 address	32	Yes for NAT44
sourceIPv6 address	128	Yes for NAT64

Table 15: Maximum fragments pending reassembly Exceeded event template

5.6.8. Threshold reached events

This event will be generated when a NAT device reaches a operator configured threshold when allocating resources. The threshold reached events are described in the section above. The following is a template of the individual events.

5.6.8.1. Address pool high or low threshold reached

This event is generated when the high or low threshold is reached for the address pool. The template is the same for both high and low threshold events

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natThresholdEvent	32	Yes
natPoolID	32	Yes
configuredLimit	32	Yes

Table 16: Address pool high/low threshold reached event template

5.6.8.2. Address and port high threshold reached

This event is generated when the high threshold is reached for the address pool and ports.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natThresholdEvent	32	Yes
configuredLimit	32	Yes

Table 17: Address port high threshold reached event template

5.6.8.3. Per-user Address and port high threshold reached

This event is generated when the high threshold is reached for the per-user address pool and ports.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natThresholdEvent	32	Yes
configuredLimit	32	Yes
vlanID/ingressVRFID	32	No
sourceIPv4 address	32	Yes for NAT44
sourceIPv6 address	128	Yes for NAT64

Table 18: Per-user Address port high threshold reached event template

5.6.8.4. Global Address mapping high threshold reached

This event is generated when the high threshold is reached for the per-user address pool and ports. This is generated only by NAT devices that use a paired address pooling behavior.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
natThresholdEvent	32	Yes
configuredLimit	32	Yes
vlanID/ingressVRFID	32	No

Table 19: Global Address mapping high threshold reached event template

5.6.9. Address binding create and delete events

These events will be generated when a NAT device binds a local address with a global address and when the global address is freed. These binding events happen when the first packet of the first flow from a host in the private realm.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
sourceIPv4 address	32	Yes for NAT44
sourceIPv6 address	128	Yes for NAT64
Translated Source IPv4 Address	32	Yes

Table 20: NAT Address Binding template

5.6.10. 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.

portRangeStart represents the starting value of the range.

portRangeEnd represents the ending value of the range.

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 and portRangeEnd MUST be specified.

It is up to the implementation to choose to consolidate log records in case two consecutive port ranges for the same user are allocated or freed.

Field Name	Size (bits)	Mandatory
timeStamp	64	Yes
natInstanceID	32	No
natEvent	8	Yes
sourceIPv4 address	32	Yes for NAT44
sourceIPv6 address	128	Yes for NAT64
Translated Source IPv4 Address	32	Yes
portRangeStart	16	Yes
portRangeEnd	16	No

Table 21: NAT Port Block Allocation event template

6. Encoding

6.1. IPFIX

This document uses IPFIX as the encoding mechanism to describe the logging of NAT events. However, the information that is 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 IEs specified for logging MUST be the same irrespective of the encoding mechanism used.

7. Acknowledgements

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

8. IANA Considerations

8.1. Information Elements

IANA will register the following IEs in the IPFIX Information Elements registry at http://www.iana.org/assignments/ipfix/ipfix.xml

8.1.1. natInstanceID

Name : natInstanceID

Description: This Information Element identifies an Instance of the NAT uniquely, that runs on a NAT middlebox function after the packet passed the Observation Point. natInstanceID is defined in RFC 7659 [RFC7659]

Abstract Data Type: unsigned32

Data Type Semantics: identifier

Reference:

See RFC 791 [RFC791] for the definition of the IPv4 source address field. See RFC 3022 [RFC3022] for the definition of NAT. See RFC 3234 [RFC3234] for the definition of middleboxes.

8.1.2. internalAddressRealm

Name: internalAddressRealm

Description: This Information Element represents the internal address realm where the packet is originated from or destined to. By definition, a NAT mapping can be created from two address realms, one from internal and one from external. Realms are implementation dependent and can represent a VRF ID or a VLAN ID or some unique identifier. Realms are optional and when left unspecified would mean that the external and internal realms are the same.

Abstract Data Type: octetArray

Data Type Semantics: identifier

Reference:

See RFC 791 [RFC791] for the definition of the IPv4 source address field. See RFC 3022 [RFC3022] for the definition of NAT. See RFC 3234 [RFC3234] for the definition of middleboxes.

8.1.3. externalAddressRealm

Name: externalAddressRealm

Description: This Information Element represents the external address realm where the packet is originated from or destined to. The detailed definition is in the internal address realm as specified above.

Abstract Data Type: octetArray

Data Type Semantics: identifier

Reference:

See RFC 791 [RFC791] for the definition of the IPv4 source address field. See RFC 3022 [RFC3022] for the definition of NAT. See RFC 3234 [RFC3234] for the definition of middleboxes.

8.1.4. natQuotaExceededEvent

Name : natQuotaExceededEvent

Description: This Information Element identifies the Quota Exceeded type that is reported by the natQuotaExceeded event. This definition shall be maintained by a new registry for values of this Information

Element, to which new values can be added. IANA maintains the registry for values of this Information Element at http://www.iana.org/assignments/ipfix/ ipfix.xml#TBD-by-IANA. New assignments for the registry will be administered by IANA, on a First Come First Served basis [RFC5226], subject to Expert Review [RFC5226]. Experts need to check definitions of new values for completeness, accuracy, and redundancy. Those experts will initially be drawn from the Working Group Chairs and document editors of the BEHAVE Working Group.

Initial values to in the registry are defined by the table below.

Quota Exceeded Event Name	Values
Maximum Session entries	1
Maximum BIB entries	2
Maximum entries per user	3
Maximum active hosts or subscribers	4
Maximum fragments pending reassembly	5

Table 22

Abstract Data Type: unsigned32

Data Type Semantics: identifier

Reference:

See RFC 791 [RFC791] for the definition of the IPv4 source address field. See RFC 3022 [RFC3022] for the definition of NAT. See RFC 3234 [RFC3234] for the definition of middleboxes.

8.1.5. natThresholdEvent

Name: natThresholdEvent

Description: This Information Element identifies the threshold type that is reported by the event. This definition shall be maintained by a new registry for values of this Information Element, to which new values can be added. IANA maintains the registry for values of this Information Element at http://www.iana.org/assignments/ipfix/ipfix.xml#TBD-by-IANA. New assignments for the registry will be administered by IANA, on a First Come First Served basis [RFC5226], subject to Expert Review [RFC5226]. Experts need to check definitions of new values for completeness, accuracy, and redundancy.

Those experts will initially be drawn from the Working Group Chairs and document editors of the BEHAVE Working Group.

Initial values to in the registry are defined by the table below.

Threshold Exceeded Event Name	Values
Address pool high threshold event	1
Address pool low threshold event	2
Address and port mapping high threshold event	3
Address and port mapping per user high threshold event	4
Global Address mapping high threshold event	5

Table 23

Abstract Data Type: unsigned32

Data Type Semantics: identifier

Reference:

See RFC 791 [RFC791] for the definition of the IPv4 source address field. See RFC 3022 [RFC3022] for the definition of NAT. See RFC 3234 [RFC3234] for the definition of middleboxes.

8.1.6. natEvent

Name : natEvent

Description: This Information Element identifies a NAT event. This IE uniquely describes an occurence of an event. Examples of NAT events include but not limited to, creation or deletion of a NAT translation entry, a threshold reached or execeeded etc. The original definitions of this Information Element specified only three values 1, 2, and 3. This definition is replaced by a new registry for values of this Information Element, to which new values can be added. The semantics of the three originally defined values remains unchanged. IANA maintains the registry for values of this Information Element at http://www.iana.org/assignments/ipfix/ipfix.xml#TBD-by-IANA. New assignments for the registry will be administered by IANA, on a First Come First Served basis [RFC5226], subject to Expert Review [RFC5226]. Experts need to check definitions of new values for completeness, accuracy, and redundancy. Those experts will initially be drawn from the Working Group Chairs and document editors of the BEHAVE Working Group.

Initial values to the registry are defined by the table below.

Event Name	Values
NAT Session create	1
NAT Session Delete	2
NAT Addresses exhausted	3
NAT44 Session create	4
NAT44 Session delete	5
NAT64 Session create	6
NAT64 Session delete	7
NAT44 Binding Information Base entry create	8
NAT44 Binding Information Base entry delete	9
NAT64 Binding Information Base entry create	10
NAT64 Binding Information Base entry delete	11
NAT ports exhausted	12
Quota exceeded	13
Address binding create	14
Address binding delete	15
Port block allocation	16
Port block de-allocation	17
Threshold reached	18

Table 24

Abstract Data Type: unsigned8

Data Type Semantics: identifier

Element ID : 230

Reference:

See RFC 3022 [RFC3022] for the definition of NAT. See RFC 3234 [RFC3234] for the definition of middleboxes. See [thisRFC] for the definitions of values 4-16.

9. Management Considerations

This section considers requirements for management of the log system to support logging of the events described above. It first covers requirements applicable to log management in general. Any additional standardization required to fullfil these requirements is out of scope of the present document. Some management considerations are covered in [I-D.behave-syslog-nat-logging]. This document covers the additional considerations.

9.1. Ability to collect events from multiple NAT devices

An IPFIX collector MUST be able to collect events from multiple NAT devices and be able to decipher events based on the Observation Domain ID in the IPFIX header.

9.2. Ability to suppress events

The exhaustion events can be overwhelming during traffic bursts and hence SHOULD be handled by the NAT devices to rate limit them before sending them to the collectors. For eg. when the port exhaustion happens during bursty conditions, instead of sending a port exhaustion event for every packet, the exhaustion events SHOULD be rate limited by the NAT device.

10. Security Considerations

The security considerations listed in detail for IPFIX in [RFC7011] applies to this draft as well. As described in [RFC7011] the messages exchanged between the NAT device and the collector MUST be protected to provide confidentiality, integrity and authenticity. Without those characteristics, the messages are subject to various kinds of attacks. These attacks are described in great detail in [RFC7011].

This document re-emphasizes the use of TLS or DTLS for exchanging the log messages between the NAT device and the collector. The log events sent in clear text can result in confidential data being exposed to attackers, who could then spoof log events based on the information in clear text messages. Hence, the log events SHOULD NOT be sent in clear text.

11. References

11.1. 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>.

[RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, DOI 10.17487/RFC2663, August 1999, <http://www.rfc-editor.org/info/rfc2663>.

[RFC4787] Audet, F., Ed. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <http://www.rfc-editor.org/info/rfc4787>.

[RFC5382] Guha, S., Ed., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, DOI 10.17487/RFC5382, October 2008, <http://www.rfc-editor.org/info/rfc5382>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <http://www.rfc-editor.org/info/rfc6146>.

[RFC6302] Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging Recommendations for Internet-Facing Servers", BCP 162, RFC 6302, DOI 10.17487/RFC6302, June 2011, <http://www.rfc-editor.org/info/rfc6302>.

[RFC6888] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, DOI 10.17487/RFC6888, April 2013, <http://www.rfc-editor.org/info/rfc6888>.

[RFC7659] Perreault, S., Tsou, T., Sivakumar, S., and T. Taylor, "Definitions of Managed Objects for Network Address Translators (NATs)", RFC 7659, DOI 10.17487/RFC7659, October 2015, <http://www.rfc-editor.org/info/rfc7659>.

11.2. Informative References

[I-D.ietf-behave-syslog-nat-logging] Chen, Z., Zhou, C., Tsou, T., and T. Taylor, "Syslog Format for NAT Logging", draft-ietf-behave-syslog-nat-logging-06 (work in progress), January 2014.

[IPFIX-IANA] IANA, "IPFIX Information Elements registry", <http://www.iana.org/assignments/ipfix>.

[RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, DOI 10.17487/RFC5470, March 2009, <http://www.rfc-editor.org/info/rfc5470>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <http://www.rfc-editor.org/info/rfc7011>.

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

Email: repenno@cisco.com