

Behave Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 25, 2014

Z. Chen
China Telecom
C. Zhou
Huawei Technologies
T. Tsou
Huawei Technologies (USA)
T. Taylor, Ed.
Huawei Technologies
September 21, 2013

Syslog Format for NAT Logging
draft-ietf-behave-syslog-nat-logging-03

Abstract

With the wide deployment of Carrier Grade NAT (CGN) devices, the logging of NAT-related events has become very important for various operational purposes. The logs may be required for troubleshooting, to identify a host that was used to launch malicious attacks, and/or for accounting purposes. This document identifies the events that need to be logged and the parameters that are required in the logs depending on the context in which the NAT is being used. It goes on to standardize formats for reporting these events and parameters using SYSLOG ([RFC 5424](#)). A companion document specifies formats for reporting the same events and parameters using IPFIX ([RFC 5101](#)). Applicability statements are provided in this document and its companion to guide operators and implementors in their choice of which technology to use for logging.

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 March 25, 2014.

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](http://trustee.ietf.org/license-info) 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	3
1.1.	Terminology	4
2.	Deployment Considerations	4
2.1.	Static and Dynamic NATs	4
2.2.	Realms and Address Pools	5
2.2.1.	Address Pools	6
2.3.	NAT Logging Requirements For Different Transition Methods	7
2.3.1.	IP Addresses and Generalized Addresses	8
2.4.	The Port Control Protocol (PCP)	9
2.5.	Logging At the Customer Edge	9
3.	NAT-Related Events and Parameters	9
3.1.	NAT Session Creation and Deletion	10
3.1.1.	Destination Logging	11
3.2.	Binding Information Base Entry Creation and Deletion . .	12
3.3.	Address Mapping Creation and Deletion Events	13
3.4.	Port Set Allocation and Deallocation	14
3.5.	Address Pool High- and Low-Water-Mark Threshold Events .	16
3.6.	Global Address Mapping High-Water-Mark Threshold Event .	17
3.7.	Global Address Mapping Limit Exceeded	17
3.8.	Global Transport Mapping High-Water-Mark Threshold Event	18
3.9.	Global Transport Mapping Limit Exceeded	18
3.10.	Subscriber-Specific Mapping Threshold Event	19
3.11.	Global Limit On Number of Active Subscribers Exceeded . .	20
3.12.	Subscriber-Specific Limit On Number of Transport Mappings Exceeded	21
3.13.	Quota Exceeded Event	21
3.14.	Global Limit On Number Of Fragments Pending Reassembly Exceeded	23
4.	SYSLOG Applicability	23
5.	SYSLOG Record Format For NAT Logging	24
5.1.	SYSLOG HEADER Fields	25

5.2.	Parameter Encodings	25
5.2.1.	APoolId: Address Pool Identifier	26
5.2.2.	DevID: Reporting Device Identifier	26
5.2.3.	DevTyp: Reporting Device Type	26
5.2.4.	PostS4: Mapped External IPv4 Address	27
5.2.5.	PostSPt: Mapped External Port or ICMP Identifier	27
5.2.6.	PreSPt: Internal Port or ICMP Identifier	27
5.2.7.	Proto: Protocol Identifier	27
5.2.8.	PScop: Protocol Scope For Quota	27
5.2.9.	PSID: Port Set Identifier	27
5.2.10.	PtRg: Allocated Port Range	27
5.2.11.	SiteID: Subscriber Site Identifier	27
5.2.12.	SScop: Site Scope For Quota	28
5.2.13.	TrigR: Realm Triggering Session Creation	28
5.2.14.	VLANid: VLAN Identifier	28
5.2.15.	VRFid: VPN Routing and Forwarding Identifier	28
5.3.	Encoding Of Complete Log Report For Each Event Type	28
5.3.1.	NAT Session Creation and Deletion	28
5.3.1.1.	Examples	29
5.3.2.	Address Binding Event	29
5.3.3.	Port Allocation Change	30
5.3.4.	Address Exhaustion Event	31
5.3.5.	NAT Port Exhaustion	31
5.3.6.	Quota Exceeded	32
6.	IANA Considerations	33
7.	Management Considerations	34
8.	Security Considerations	34
9.	References	34
9.1.	Normative References	35
9.2.	Informative References	35
	Authors' Addresses	37

[1.](#) Introduction

Operators already need to record the addresses assigned to subscribers at any point in time, for operational and regulatory reasons. When operators introduce NAT devices which support address sharing (e.g., Carrier Grade NATs (CGNs)) into their network, additional information has to be logged. This document and [\[I-D.behave-ipfix-nat-logging\]](#) are provided in order to standardize the events and parameters to be recorded, using SYSLOG [\[RFC5424\]](#) and IPFIX [\[RFC5101\]](#) respectively. The content proposed to be logged by the two documents is exactly the same, but as will be seen, the choice of which to use in a given scenario is an engineering issue.

Detailed logging requirements will vary depending on the context in which they are used. For example, different methods for transition from IPv4 to IPv6 require different events and different parameters to be logged. [Section 2](#) covers this topic.

[Section 3](#) provides a more detailed description of the events that need logging and the parameters that may be required in the logs.

The use of SYSLOG [[RFC5424](#)] has advantages and disadvantages compared with the use of IPFIX [[RFC5101](#)]. [Section 4](#) provides a statement of applicability for the SYSLOG approach.

[Section 5](#) specifies SYSLOG record formats for logging of the events and parameters described in [Section 3](#). The definitions provide the flexibility to vary actual log contents based on the requirements of the particular deployment.

[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 "Key words for use in RFCs to Indicate Requirement Levels" [[RFC2119](#)].

This document uses the terms "session" and Binding Information Base (BIB) as they are defined in [Section 2 of \[RFC6146\]](#). Note that this definition of "session" is destination-specific, where the original definition of a NAT session in [[RFC2663](#)] is destination-independent.

This document uses the term "address mapping" to denote the initial logical step required to set up a session, as described in [Section 2.2](#). It uses the term "transport binding" to denote the content of a BIB entry.

Except where a clear distinction is necessary, this document uses the abbreviation "NAT" to encompass both Network Address Translation (NAT in the strict sense) and Network Address and Port Translation (NAPT). The event report descriptions provided in this document apply to NAPT, and can be simplified for pure NAT operation.

[2.](#) Deployment Considerations

[2.1.](#) Static and Dynamic NATs

A NAT controls a set of resources in the form of one or more pools of external addresses. If the NAT also does port translation (i.e., it is a NAPT), it also controls the sets of UDP and TCP port numbers and ICMP identifiers associated with each external address.

Logging requirements for a NAT depend heavily on its resource allocation strategy. NATs can be classed as static or dynamic depending on whether the resources provided to individual users are pre-configured or allocated in real time as the NAT recognizes new flows.

Static assignments can be logged at configuration time by the NAT or by network infrastructure. The logging volume associated with static assignments will be relatively low, of the order of the volume of user logons. As discussed below, static assignments are typically associated with IPv6 transition methods rather than traditional NAT. The details of what to log will depend on the transition method concerned.

Dynamic assignments typically require both more detail in the logs and a higher volume of logs in total. A traditional Network Address Port Translator (NAPT) as described in [\[RFC3022\]](#) and following the recommendations of [\[RFC4787\]](#) and [\[RFC5382\]](#) will generate a new mapping each time it encounters a new internal <address, port> combination.

For statistical reasons, static assignments support lower address sharing ratios than fully dynamic assignments as exemplified by the traditional NAPT. The sharing ratio can be increased while restraining log volumes by assigning ports to users in multi-port increments as required rather than assigning just one port at a time. A subscriber may start with no initial allocation, or may start with an initial permanent allocation to which temporary increments are added when the initial set is all being used. See [\[RFC6264\]](#) and [\[I-D.tsou-behave-natx4-log-reduction\]](#) for details. If this strategy is followed, logging will be required only when an increment is allocated or reclaimed rather than every time an internal <address, port> combination is mapped to an external <address, port>.

2.2. Realms and Address Pools

A realm defines the scope within which a specific set of addresses are unique. In general these will be IPv4 or IPv6 addresses, but not necessarily. A counter-example specifically addressed by this document is the case of Gateway-Initiated DS-Lite [\[RFC6674\]](#), where individual host sites are identified by 16-bit Softwire Identifiers. See further discussion in [Section 2.3](#) and [Section 2.3.1](#).

Table [proposed] in the NAT-MIB [\[I-D.Behave-NAT-MIB\]](#) provides a mapping between each realm identifier and the Virtual Routing Function (VRF) instance, VLAN identifier, or Gateway-Initiated DS-Lite context, if any, with which it is associated.

From the point of view of a specific NAT session, only two realms are involved: an internal realm and an external realm. However, the NAT as a whole may support a number of realms, for example:

- o multiple internal realms with overlapping address spaces;
- o an external IPv4 public realm; and/or
- o an external IPv6 public realm.

As described in [[RFC6146](#)], for example, setting up a NAT session proceeds in a series of logical steps. The first step in particular may not be implemented explicitly in a given implementation, but logically it has to happen before the next step can be taken.

1. An address mapping is created between the internal realm and an external realm chosen based on information in the triggering packet or administrative request.
2. Using that address mapping, a transport binding is created between specific transport endpoints (e.g., between specific port values) in the two realms for the protocol required by the session, and added to the Binding Information Base (BIB).
3. Setup of the session is completed by mapping the destination address and port (if necessary) into the selected external realm.

This section is concerned only with the address mapping step. That step is always triggered either by a packet outgoing from the internal host to a given destination, or by administrative action providing equivalent information. The external realm for the mapping is chosen based on the destination.

To summarize where we are: an address mapping binds an internal address with an external address in a selected external realm. One address mapping can serve as the basis for one to many transport bindings in the BIB, and one BIB entry can serve as the basis for one to many sessions. A single internal address may be associated with multiple address mappings at one time.

[2.2.1.](#) Address Pools

An address pool is a mechanism for configuring the set of addresses to which a given internal address can be mapped in a given realm. The pool may be used simply to ration the available addresses within that realm, or may be selected for other reasons such as to add additional semantics (e.g., type of service required) to the external address within the target realm. Clearly a given internal address may be mapped into more than one address pool at a given time.

The model of an address pool assumed in this document and in the NAT MIB [[I-D.Behave-NAT-MIB](#)] is that the pool offers a fixed range of port/ICMP identifier values, the same over all addresses within the pool. How these are allocated to individual transport bindings in the BIB depends on the pooling behaviour. With a pooling behaviour of "arbitrary" [[RFC4787](#)], the NAT can select any address in the pool with a free port value for the required protocol and map the internal address to it. With the recommended pooling behaviour of "paired" [[RFC4787](#)], the NAT restricts itself to finding a free port at the address to which the internal address is already mapped, if there is one.

From this description, one can see that ports are a limited resource, subject to exhaustion at the pool level and, with "paired" behaviour, at the level of the individual address. Log events are defined in [Section 3.5](#) that allow monitoring of port utilization at the pool level. [Section 7](#) discusses how the thresholds for triggering these events should be varied depending on pooling behaviour.

2.3. NAT Logging Requirements For Different Transition Methods

A number of transition technologies have been or are being developed to aid in the transition from IPv4 to IPv6. 6rd [[RFC5969](#)] and DS-Lite [[RFC6333](#)] are at the deployment stage. Several 'stateless' technologies: Public IPv4 over IPv6 [[I-D.softwire-public-4over6](#)], MAP-E [[I-D.softwire-map](#)], and Lightweight 4over6 [[I-D.softwire-lw4over6](#)] have seen experimental deployment and are in the process of being standardized at the time of writing of this document.

Of the technologies just listed, 6rd and Public IPv4 over IPv6 do not involve NATs and hence need not be considered further. The other techniques involve NAT at the customer edge, at the border router, or both, and hence are in scope.

A DS-Lite Address Family Transition Router (AFTR) includes a large-scale session-stateful NAT44 processing potentially millions of sessions per second. The special character of AFTR operation over that of a traditional NAT44 is that the source IPv4 addresses of the internal hosts may not be unique. As a consequence, the session

tables need to include an alternative identifier associated with the subscriber host. For basic DS-Lite, this will be the IPv6 address used to encapsulate the packets outgoing from the host. See [Section 6.6 of \[RFC6333\]](#). For gateway-initiated DS-Lite [\[RFC6674\]](#), an identifier associated with the incoming tunnel from the host is used instead.

The DS-Lite customer edge equipment (the 'B4') may also perform NAT44 functions, similar to the functions performed by traditional NAT44 devices.

As a NAT44, the DS-Lite AFTR may be fully dynamic, or may allocate ports in increments as described in the previous section.

Lightweight 4over6 [\[I-D.softwire-lw4over6\]](#) and MAP-E [\[I-D.softwire-map\]](#) both require NAT44 operation at the customer equipment (unified CPE, [\[I-D.softwire-unified-cpe\]](#)). In both cases the resource allocation strategy is static. Thus any logging of resource allocation for these two transition techniques can be done by the network at configuration time.

2.3.1. IP Addresses and Generalized Addresses

In the event reports described below, external addresses and destination addresses will always be true IPv4 or IPv6 addresses. Source addresses of outgoing packets before mapping will also be IP addresses, but will not always be meaningful because they will not be unique within their realm. This is true in particular of some of the transition methods described in the previous section.

For this reason, the event report descriptions introduce the term "generalized address" to describe internal addresses (as opposed to source addresses within packets). The detailed description of the encoding of a generalized address in [Section 5.2](#) provides for an address type, address/prefix value, and address/prefix length, similarly to the encoding of an IP address. However, the range of generalized address types is expanded to support the following:

- o For traditional NATs, the source IPv4 address (for NAT44) or IPv6 address (for NAT64) is sufficient.
- o For the DS-Lite, Lightweight 4over6 or MAP-E transition methods, the subscriber site can be identified by the IPv6 tunnel endpoint prefix or address provisioned to that site.
- o Gateway-initiated DS-Lite uses the combination of a 32-bit context identifier (CID) and a softwire identifier (SWID). Several different realizations of these identifiers are described in

[Section 6 of \[RFC6674\]](#). From the point of view of this document, the CID is represented by a realm identifier, leaving the SWID as the value of the generalized address itself.

2.4. The Port Control Protocol (PCP)

The Port Control Protocol (PCP) [[RFC6887](#)] and its port set extension [[I-D.pcp-port-set](#)] can be viewed as a way to provision ports by other means. However, PCP can be invoked on a per-flow basis, so the volume of logs generated by a PCP server can be closer to the volume associated with a fully dynamic NAT. The volume really depends on how PCP is being used in a specific network.

2.5. Logging At the Customer Edge

Logging at the customer edge (or at the ISP edge for NATs protecting the ISP's internal networks) may be done by the customer for purposes of internal management, or by the ISP for its own administrative and regulatory purposes. Given the likelihood of a high internal community of interest, it is possible but unlikely that a NAT at the edge of a large enterprise network processes a number of new packet flows per second which is comparable to the volume handled by a carrier grade NAT. Most customer edge NATs will handle a much smaller volume of flows.

3. NAT-Related Events and Parameters

The events which follow were initially gleaned, in the words of the authors of [[I-D.behave-ipfix-nat-logging](#)], from [[RFC4787](#)] and [[RFC5382](#)]. Some details were subsequently informed by the discussion in [Section 2](#) and by provisions within the NAT MIB [[I-D.Behave-NAT-MIB](#)]. [Section 4 of \[RFC6888\]](#) also provides a brief statement of logging requirements for carrier grade NATs.

Since the present document deals with SYSLOG rather than IPFIX, the timestamp and the event type will appear in the log header rather than as an explicit part of the structured data portion of the log. Hence they are omitted from the parameter tabulations that follow.

The listed parameters include an optional triggering NAT procedure in each case. The triggering NAT procedure is an implementation-defined string indicating the type of NAT function (e.g., NAT44) being applied in association with the reported event. The same device can offer different functions depending on the particular packets being processed. The triggering NAT procedure is most relevant when address mappings are created, but if address mapping events are not reported (i.e., because pooling behaviour is "arbitrary" [[RFC4787](#)]), implementations MAY choose to report it for BIB entry or session events.

3.1. NAT Session Creation and Deletion

A NAT session creation or deletion event is logged when a transport binding is further bound to or unbound from a specific destination address and port in the external realm. One to many sessions can be based on the same transport binding.

Implementations MUST NOT report session creation and deletion events unless destination logging is enabled (see discussion below).

The following specific events are defined:

- o NAT session creation
- o NAT session deletion

These take the same parameters for all types of NAT. Parameters "internal realm" through "protocol identifier" capture the underlying transport binding. The destination IP address and port and possibly the trigger are unique to the session. If the destination IP and port do not require remapping into the external realm, the internal values are redundant and SHOULD be omitted from the report. So long as the underlying BIB entry exists, the internal values can in any event be retrieved from the natMappingTable in the NAT MIB [[I-D.Behave-NAT-MIB](#)] using the combination of protocol, external realm, external destination address, and external destination port as key.

- o Triggering NAT procedure (OPTIONAL);
- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);
- o Internal port or ICMP identifier (MANDATORY);
- o External realm (MANDATORY);

- o External IP address (MANDATORY);
- o External port or ICMP identifier (MANDATORY);
- o Protocol identifier (MANDATORY);
- o Internal destination IP address (as given in outgoing packets) (OPTIONAL);
- o Internal destination port or ICMP identifier (as given in outgoing packets) (OPTIONAL);
- o External destination IP address (as given in outgoing packets) (MANDATORY);
- o External destination port or ICMP identifier (as given in outgoing packets) (MANDATORY);
- o Trigger for session creation or deletion (OPTIONAL):
 - * outgoing packet received;
 - * incoming packet received;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]); or
 - * deletion of the underlying BIB entry.

3.1.1. Destination Logging

The logging of destination address and port is generally undesirable, for several reasons. [[RFC6888](#)] recommends against destination logging because of the privacy issues it creates. From an operator's point of view, destination logging is costly not just because of the volume of logs it will generate, but because the NAT now has to carry additional session state so that it only needs to log once per session between two transport end points rather than logging every packet. Finally, [[RFC4787](#)], etc. recommend the use of endpoint-independent mapping to maximize the ability of applications to operate through the NAT. In that case, most of the contents of the session creation event report will be repeated for one destination after another.

One possibility is that the implementation provides the operator with the ability to log destinations only for particular subscribers or particular mapped addresses on a special study basis. This facility could be used for trouble-shooting or malicious activity tracing in

particular cases as required. If such a capability is provided, the implementation **MUST** report session creation and deletion events for sessions matching the specified criteria, but **MUST NOT** report these events for other sessions.

3.2. Binding Information Base Entry Creation and Deletion

A transport mapping as recorded in the Binding Information Base (BIB) corresponds to the older definition of NAT session as defined in [Section 2.3 of \[RFC2663\]](#). The BIB entry creation or deletion event reports the addition or deletion of a mapping between an internal transport endpoint and an external transport address. The event report provides the same information as the session creation/deletion event, except for the destination-related fields in the latter.

Particularly with endpoint-independent mapping behaviour [[RFC4787](#)], one BIB entry creation event is associated with potentially many succeeding session creation events, as individual destinations are mapped into the session table. Similarly, a BIB entry deletion event will be associated with potentially many session deletion events, which may have preceded it over a period of time or may occur at the same time as a result of the BIB entry deletion.

Operators **SHOULD** disable the reporting of BIB entry creation and deletion events when destination logging is enabled, because of the redundancy between the BIB and session event reports. However, in the case of endpoint-independent mapping behaviour [[RFC4787](#)], the BIB event provides a compact summary of most of the content of what could be a large number of corresponding session events.

The following specific events are defined:

- o BIB entry creation
- o BIB entry deletion

These take the same parameters for all types of NAT. The internal realm, generalized internal address, external realm, and external address capture the underlying address mapping. The port values, protocol, and possibly the trigger are unique to the BIB entry.

- o Triggering NAT procedure (OPTIONAL);
- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);
- o Internal port or ICMP identifier (MANDATORY);

- o External realm (MANDATORY);
- o External address (MANDATORY);
- o External port or ICMP identifier (MANDATORY);
- o Protocol identifier (MANDATORY);
- o Trigger for transport mapping creation or deletion (OPTIONAL):
 - * outgoing packet received;
 - * incoming packet received;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]); or
 - * deletion of the underlying address mapping.

3.3. Address Mapping Creation and Deletion Events

Two specific events are provided:

- o Address mapping creation;
- o Address mapping deletion.

Address mapping is discussed in detail in [Section 2.2](#).

One address mapping creation event is associated with potentially many succeeding BIB entry creation events, as individual port values are mapped into the BIB for specific protocols. Similarly, an address mapping deletion event will be associated with potentially many BIB entry deletion events, which may have preceded it over a period of time or may occur at the same time as a result of the address unbinding.

The address mapping events take the following specific parameters:

- o Triggering NAT procedure (OPTIONAL);
- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);
- o External realm (MANDATORY);
- o External IP address (MANDATORY).

- o Trigger for address mapping creation or deletion (OPTIONAL):
 - * outgoing packet;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]); or
 - * autonomous action of the NAT.

3.4. Port Set Allocation and Deallocation

This event is recorded at a hybrid NAT whenever the set of ports allocated to a given address mapping changes. It is assumed that when ports are allocated in bulk, the same values are allocated for all protocols.

The following specific events are defined:

- o Port set allocation;
- o Port set deallocation.

The parameters for these events are:

- o Triggering NAT procedure (OPTIONAL);
- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);
- o External realm (MANDATORY);
- o External IP address (MANDATORY);
- o A set of ports available for transport mapping, newly allocated to or deallocated from the given address mapping. The representation of a port set is described in the next paragraph (MANDATORY).

A port set is represented by four parameters. The full set of parameters describes a sequence of equally-spaced and equally-sized ranges of consecutive port values. If only a single range is allocated or deallocated, two of the parameters can be omitted. The four parameters are:

- o Starting port number, the lowest port number in the entire port set (MANDATORY);

- o Ending port number, the highest port number in the entire port set (MANDATORY);
- o Range size, the number of port values in each range (OPTIONAL);
- o Range step, the difference between the first port number in one range and the first port number in the immediately preceding range of the port set (OPTIONAL).

In the case of a single range, range size SHOULD be omitted and range step MUST be omitted because it is meaningless.

Examples:

1. Two ranges, 1024-1535 and 2048-2559 are allocated. Each range consists of 512 consecutive port numbers. The parameter values to represent this allocation are:
 - * starting port = 1024
 - * ending port = 2559
 - * range size = 512
 - * range step = 1024.
2. Strictly for purposes of illustration, assume a sequence of 512 even-numbered ports is allocated, beginning at 1024, then 1026, ending at 2046. The parameter values to represent this allocation are:
 - * starting port = 1024
 - * ending port = 2046
 - * range size = 1
 - * range step = 2.
3. A single range of ports is allocated, running consecutively from 1024 to 2046. The parameter values to represent this allocation are:
 - * starting port = 1024
 - * ending port = 2046.

It will be necessary to use multiple event reports to report more complex allocations or deallocations.

3.5. Address Pool High- and Low-Water-Mark Threshold Events

Two specific events provide reports on address pool utilization:

- o High-water-mark threshold reached or over-shot;
- o Low-water-mark threshold reached or under-shot.

Depending on deployment the operator has the alternative of using the SNMP notifications `natNotifPoolWater-MarkHigh` and `natNotifPoolWater-MarkLow` defined in the NAT MIB [[I-D.Behave-NAT-MIB](#)] rather than logging these events.

Address pools are discussed in [Section 2.2.1](#). The `natPoolTable` object in the NAT MIB [[I-D.Behave-NAT-MIB](#)] provides access to parameters describing the utilization level of address-port combinations within a given pool. Since a new transport mapping cannot be allocated unless a mappable address and a free port on that address are available, it is important to know when the available set of address-port combinations within a given pool is nearing exhaustion. Hence the `natPoolTable` contains a high-water-mark threshold settable by the operator. An address pool high-water-mark event report is generated when a new mapping into the pool is triggered and aggregate address-port utilization is equal to or greater the threshold.

Similarly it can be of interest to know when a pool is under-utilized. Hence the `natPoolTable` also provides a low-water-mark threshold. An address pool low-water-mark event report is generated when aggregate address-port utilization is equal to or less than the low-water-mark threshold.

[Section 7](#) discusses factors affecting the choice of the threshold values, taking note that the port utilization as computed does not take account of the number of different protocols mapped to a given port value.

An address pool threshold event report contains the following specific parameters:

- o Triggering NAT procedure (OPTIONAL);
- o Pool identifier, equal to the value of the `natPoolIndex` object presented in the `natPoolTable` in the MIB (MANDATORY).

3.6. Global Address Mapping High-Water-Mark Threshold Event

One specific event allows monitoring of the total number of mappings between internal and external addresses:

- o Address mapping high-water-mark threshold equalled or exceeded.

This event report is most meaningful when the pooling type behaviour is "paired" [[RFC4787](#)], and is especially applicable to devices implementing NAT functionality only and not port translation. Depending on deployment, operators can choose instead to use the SNMP notification `natNotifAddrMappings` defined in the NAT MIB [[I-D.Behave-NAT-MIB](#)].

The NAT MIB displays cumulative counts of address mappings created and removed in the `natCounters` table. When the difference between these two counters equals or exceeds the threshold `natAddrMapNotifyThreshold` provided in the `natLimits` table the global address binding high-water-mark threshold event is reported.

The specific parameters provided by this event report are:

- o Triggering NAT procedure (OPTIONAL);
- o Current number of active address mappings, equal to the difference between the `natAddressMappingCreations` and `natAddressMappingRemovals` counters displayed in the `natCounters` table in the NAT MIB (MANDATORY).

3.7. Global Address Mapping Limit Exceeded

The "Global Address Mapping Limit Exceeded" event is reported when a new address mapping is triggered but the total number of address mappings would exceed an administrative limit if it were added. The limit is given by object `natLimitAddressMappings` in the `natLimits` table of the NAT MIB. MIB counters giving number of packets dropped due to resource limitations including this one are:

- o globally, `natResourceErrors` in the `natCounters` table;
- o per protocol, `natProtocolResourceErrors` in `natProtocolTable`;
- o per subscriber, `natSubscriberResourceErrors` in `natSubscribersTable`.

The parameters for this event are:

- o Triggering NAT procedure (OPTIONAL);

- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);
- o Trigger for address mapping creation (MANDATORY):
 - * outgoing packet;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]).

[3.8.](#) Global Transport Mapping High-Water-Mark Threshold Event

One specific event allows monitoring of the total number of mapping entries in the Binding Information Base (BIB):

- o Transport mapping high-water-mark threshold equalled or exceeded.

Depending on deployment, operators can choose instead to use the SNMP notification `natNotifMappings` defined in the NAT MIB [[I-D.Behave-NAT-MIB](#)].

The NAT MIB displays cumulative counts of mappings created in and removed from the BIB in the `natCounters` table. When the difference between these two counters equals or exceeds the threshold `natMappingsNotifyThreshold` provided in the `natLimits` table the global mapping high-water-mark threshold event is reported.

The specific parameters provided by this event report are:

- o Reporting device identifier (OPTIONAL);
- o Triggering NAT procedure (OPTIONAL);
- o Current number of active mappings, equal to the difference between the `natMappingCreations` and `natMappingRemovals` counters displayed in the `natCounters` table in the NAT MIB (MANDATORY).

[3.9.](#) Global Transport Mapping Limit Exceeded

The "Global Transport Mapping Limit Exceeded" event is reported when a new transport mapping (i.e., BIB entry creation) is triggered but the total number of transport mappings would exceed an administrative limit if it were added. The limit is given by object `natLimitMappings` in the `natLimits` table of the NAT MIB. MIB counters giving number of packets dropped due to resource limitations including this one are:

- o globally, natResourceErrors in the natCounters table;
- o per protocol, natProtocolResourceErrors in natProtocolTable;
- o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event are:

- o Reporting device identifier (OPTIONAL);
- o Triggering NAT procedure (OPTIONAL);
- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);
- o Trigger for BIB entry creation (MANDATORY):
 - * incoming packet;
 - * outgoing packet;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]).

3.10. Subscriber-Specific Mapping Threshold Event

An event is provided to allow monitoring of the total number of BIB entries per subscriber:

- o Subscriber-specific mapping high-water-mark threshold equalled or exceeded.

Depending on deployment, operators can choose instead to use the SNMP notification natNotifSubscriberMappings defined in the NAT MIB [[I-D.Behave-NAT-MIB](#)].

The NAT MIB displays cumulative counts of transport mappings created and removed per subscriber in the natSubscribersTable. When the difference between these two counters equals or exceeds the threshold natSubscriberMapNotifyThresh provided in that table the subscriber mapping high-water-mark threshold event is reported.

The specific parameters provided by this event report are:

- o Reporting device identifier (OPTIONAL);

- o Triggering NAT procedure (OPTIONAL);
- o Internal realm of the subscriber (MANDATORY);
- o Generalized internal address of the subscriber (MANDATORY);
- o Current number of active transport mappings for this subscriber, equal to the difference between the natSubscriberMappingCreations and natSubscriberMappingRemovals counters displayed in the natSubscribersTable table in the NAT MIB (MANDATORY).

3.11. Global Limit On Number of Active Subscribers Exceeded

The "Global Limit On Number of Active Subscribers Exceeded" event is reported when an address mapping is triggered (at least at the logical level) for a subscriber with no previous active mappings, but the total number of active subscribers would exceed an administrative limit if it were added. The limit is given by object natLimitSubscribers in the natLimits table of the NAT MIB. MIB counters giving number of packets dropped due to resource limitations including this one are:

- o globally, natResourceErrors in the natCounters table;
- o per protocol, natProtocolResourceErrors in natProtocolTable;
- o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event are:

- o Reporting device identifier (OPTIONAL);
- o Triggering NAT procedure (OPTIONAL);
- o Internal realm of the rejected subscriber (MANDATORY);
- o Generalized internal address of the rejected subscriber (MANDATORY);
- o Trigger for mapping creation (MANDATORY):
 - * outgoing packet;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]).

3.12. Subscriber-Specific Limit On Number of Transport Mappings Exceeded

The "Subscriber-Specific Limit On Number of Transport Mappings Exceeded" event is reported when a new BIB entry is triggered, but the total number of BIB entries for that subscriber would exceed an administrative limit if it were added. The limit is given by object `natSubscriberLimitMappings` in `natSubscribersTable` in the NAT MIB. MIB counters giving number of packets dropped due to resource limitations including this one are:

- o globally, `natResourceErrors` in the `natCounters` table;
- o per protocol, `natProtocolResourceErrors` in `natProtocolTable`;
- o per subscriber, `natSubscriberResourceErrors` in `natSubscribersTable`.

The parameters for this event are:

- o Triggering NAT procedure (OPTIONAL);
- o Internal realm of the subscriber (MANDATORY);
- o Generalized internal address of the subscriber (MANDATORY);
- o Trigger for transport mapping creation (MANDATORY):
 - * incoming packet;
 - * outgoing packet;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]).

3.13. Quota Exceeded Event

A "Quota Exceeded" event is reported when the NAT cannot allocate a new address mapping, transport mapping, or session because an administrative quota has been reached. Quotas may be applied on absolute quantities or on rates. The specific types of quota capability offered by a device are implementation dependent, hence the "Quota Exceeded" event reports only the minimum of information needed to identify and interpret the quota. Table [proposed] in the NAT MIB lists quota identifiers and corresponding total counts of packets dropped because of quota violations. This table may be extended to provide information on the configuration of the particular quota, depending on the implementation.

A number of counters within the NAT MIB record the number of packets dropped due to quota violations:

- o globally, in counter natQuotaDrops in the natCounters table;
- o by protocol, in the natProtocolQuotaDrops counter in the natProtocolTable;
- o per subscriber, in counter natSubscriberQuotaDrops in the natSubscribersTable.

In the list of report parameters that follows, the internal realm and generalized internal address **MUST** be provided if they are available. If the trigger for the quota violation is a packet, the contents of the received packet header and the realm that the packet came from **MUST** be reported. If the trigger was an administrative action, the equivalent to as much of this information as possible **SHOULD** be reported.

- o Triggering NAT procedure (OPTIONAL);
- o Quota identifier (MANDATORY);
- o Internal realm (OPTIONAL);
- o Generalized internal address (OPTIONAL);
- o Source realm for triggering packet (OPTIONAL);
- o Source IP address (OPTIONAL);
- o Source port or ICMP identifier (OPTIONAL);
- o Destination IP address (OPTIONAL);
- o Destination port (OPTIONAL);
- o Protocol (OPTIONAL);
- o Trigger for quota violation (OPTIONAL)
 - * packet received at the NAT;
 - * administrative action (e.g., via the Port Control Protocol [[RFC6887](#)]).

In the special case where the quota addresses bulk port allocation, the parameters listed above **MUST** be interpreted and populated as

follows, so as to capture the address mapping to which the ports would have been allocated:

- o Internal realm and generalized internal address retain their usual meanings;
- o Source realm and source IP address present the external realm and address portion of the address mapping;
- o port numbers, protocol, and destination address MUST be omitted.

3.14. Global Limit On Number Of Fragments Pending Reassembly Exceeded

The "Global Limit On Number Of Fragments Pending Reassembly Exceeded" event is reported when a new fragment is received and the number of fragments currently awaiting reassembly is already equal to an administrative limit. That limit is given by the natLimitFragments object in the natLimits table. This event MUST NOT be reported unless the NAT supports the "Receive Fragments Out of Order" behavior [[RFC4787](#)]. MIB counters giving number of packets dropped due to resource limitations including this one are:

- o globally, natResourceErrors in the natCounters table;
- o per protocol, natProtocolResourceErrors in natProtocolTable;
- o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event provide the contents of the IP header of the received fragment that triggered it. If the source realm is internal and the generalized internal address is available, it MUST also be included.

- o Source realm of the packet (MANDATORY);
- o Source IP address (MANDATORY);
- o Destination IP address (MANDATORY);
- o Generalized internal address of the source (OPTIONAL).

4. SYSLOG Applicability

The primary advantage of SYSLOG is the human readability and searchability of its contents. In addition, it has built-in priority and severity fields that allow for separate routing of reports requiring management action. Finally, it has a well-developed underpinning of transport and security protocol infrastructure.

SYSLOG presents two obstacles to scalability: the fact that the records will typically be larger than records based on a binary protocol such as IPFIX, and, depending on the architectural context, the reduced performance of a router that is forced to do text manipulation in the data plane. One has to conclude that for larger message volumes, IPFIX should be preferred as the reporting medium on the NAT itself. It is possible that SYSLOG could be used as a back-end format on an off-board device processing IPFIX records in real time, but this would give a limited boost to scalability. One concern expressed in list discussion is that when the SYSLOG formatting process gets overloaded records will be lost.

As a result, the key question is what the practical cutoff point is for the expected volume of SYSLOG records, on-board or off-board the NAT. This obviously depends on the computing power of the formatting platform, and also on the record lengths being generated.

Information has been provided to the BEHAVE list at the time of writing to the effect that one production application is generating an average of 150,000 call detail records per second, varying in length from 500 to 1500 bytes. Capacities several times this level have been reported involving shorter records, but this particular application has chosen to limit the average in order to handle peaks.

As illustrated by the examples in [Section 5.3](#), typical record sizes for the high-volume logs are in the order of 150 to 200 bytes, so throughput capacity should be higher than in the call detail case for the same amount of computing power. In private communication, a discussant has noted a practical limit of a few hundred thousand SYSLOG records per second on a router.

5. SYSLOG Record Format For NAT Logging

This section describes the SYSLOG record format for NAT logging in terms of the field names used in [\[RFC5424\]](#) and specified in [Section 6](#) of that document. In particular, this section specifies values for the APP-NAME and MSGID fields in the record header, the SD-ID identifying the STRUCTURED-DATA section, and the PARAM-NAMES and PARAM-VALUE types for the individual possible parameters within that section. The specification is in three parts, covering the header, encoding of the individual parameters, and encoding of the complete log record for each event type.

5.1. SYSLOG HEADER Fields

Within the HEADER portion of the SYSLOG record, the priority (PRI) level is subject to local policy, but a default value of 8x is suggested, representing a Facility value of 10 (security/authorization) and a Severity level varying with the event type. The suggested value by event type is shown in Table 1. The HOSTNAME field MUST identify the NAT. The value of the HOSTNAME field is subject to the preferences given in [Section 6.2.4 of \[RFC5424\]](#).

The values of the APP-NAME and MSGID fields in the record header determine the semantics of the record. The APP-NAME value "NAT" indicates that the record relates to an event reported by a NAT device. The MSGID values indicate the individual events. They are listed in Table 1 for each of the events defined in [Section 3](#). The table also shows the SD-ID value used to label the event-specific STRUCTURED-DATA element.

Event	MSGID	PRI	SD-ID
NAT session creation	SessAdd	86 info	NATsess
NAT session deletion	SessDel	86 info	NATsess
Address binding event	AddrBind	86 info	NATBind
Port allocation change	PtAlloc	86 info	NATPBlk
NAT address exhaustion	AddrEx	82 critical	NATAddrEx
NAT port exhaustion	PortEx	84 warning	NATPEX
Quota exceeded	Quota	85 notice	NATQEx
Invalid port detected	InvPort	83 error	NATInvP

Table 1: Recommended MSGID Encodings and Default PRI Values for the Events Defined In [Section 3](#)

5.2. Parameter Encodings

This section describes how to encode the individual parameters that can appear in NAT-related logs. The parameters are taken from the event descriptions in [Section 3](#), and are listed in Table 2. Formally, as will be seen in Table 9, a parameter used with more than one event is registered as multiple separate parameters, one for each event report in which it is used. However, there is no reason to change either the PARAM-NAME or the encoding of the PARAM-VALUE between different instances of the same parameter.

PARAM-NAME	Parameter
------------	-----------

depends on the reporting device type of the reporting NAT are noted in [Section 5.3](#).

[5.2.4](#). PostS4: Mapped External IPv4 Address

PARAM-VALUE: IPv4 address, represented in dotted decimal form.

[5.2.5](#). PostSPt: Mapped External Port or ICMP Identifier

PARAM-Value: decimal integer, port number or ICMP query identifier.

[5.2.6](#). PreSPt: Internal Port or ICMP Identifier

PARAM-Value: decimal integer, port number or ICMP query identifier.

[5.2.7](#). Proto: Protocol Identifier

PARAM-VALUE: an integer indicating the value of the Protocol header field (IPv4) or Next Header field (IPv6) in the incoming packet(s) (after decapsulation, for reporting device type "AFTR") to which the event described by this record applies.

[5.2.8](#). PScop: Protocol Scope For Quota

PARAM-VALUE: as for Proto for a specific protocol. "*" for sum over all protocols.

[5.2.9](#). PSID: Port Set Identifier

PARAM-VALUE: integer between 0 and 65535 designating a port set. In practice the upper limit is likely to be two orders of magnitude smaller.

[5.2.10](#). PtRg: Allocated Port Range

PARAM-VALUE: a field consisting of two decimal integers separated by a minus sign/hyphen. The first integer is the lowest port number, the second, the highest port number, in a range of consecutive ports.

[5.2.11](#). SiteID: Subscriber Site Identifier

A human-readable US-ASCII string identifying a specific host or CPE served by the reporting device. The type of identifier depends on the configuration of the reporting device, and is implementation and deployment-specific. See [Section 3](#) for a discussion of the possible identifier types.

5.2.12. SScop: Site Scope For Quota

PARAM-VALUE: "S" for single site, "M" for sum over multiple sites, served by the same VLAN or VRF. "*" for sum over all sites served by the NAT.

5.2.13. TrigR: Realm Triggering Session Creation

PARAM-VALUE: "I" for internal, "E" for external.

5.2.14. VLANid: VLAN Identifier

PARAM-VALUE: a decimal integer representing the VLAN identifier associated with the subscriber site.

5.2.15. VRFid: VPN Routing and Forwarding Identifier

PARAM-VALUE: a hexadecimal number representing a VPN identifier [RFC2685] associated with the subscriber site. It is RECOMMENDED that implementations be configurable to include or not include the OUI portion of the identifier.

5.3. Encoding Of Complete Log Report For Each Event Type

This section describes the complete NAT-related contents of the logs used to report the events listed in Table 1.

5.3.1. NAT Session Creation and Deletion

As shown in Table 1, the NAT session creation event is indicated by MSG-ID set to "SessAdd". Similarly, the NAT session deletion event is indicated by MSG-ID set to "SessDel". For both events, the associated SD-ELEMENT is tagged by SD-ID "NATsess". The contents of the NATsess SD-ELEMENT are shown in Table 3. The requirements for these contents are derived from the description in [Section 3.1](#).

PARAM-NAME	Description	Requirement
DevTyp	Section 5.2.3	OPTIONAL
DevID	Section 5.2.2	OPTIONAL
SiteID	Section 5.2.11	MANDATORY
PostS4	Section 5.2.4	MANDATORY
Proto	Section 5.2.7	MANDATORY
PreSPt	Section 5.2.6	MANDATORY
PostSPt	Section 5.2.5	MANDATORY
TrigR	Section 5.2.13	OPTIONAL

Table 3: Contents Of the SD-ELEMENT Section For Logging the Session Creation and Deletion Events

5.3.1.1. Examples

The first example is deliberately chosen to show how long a complete session log might be. For this first example, assume the log is formatted at an off-board device, which collects the information from an AFTR. Thus HOSTNAME and DevID are both present. IPv6 addresses are reported omitting a common /16 prefix and the IID portion of the address (not to be too unrealistic!). All the optional parameters are present. Note that the log could also include other SD-ELEMENTs (e.g., timeQuality), but enough is enough.

The log appears as a single record, but is wrapped between lines for purposes of presentation.

```
<86>1 2013-05-07T22:14:15.03Z record.example.net NAT 5063 SessAdd
[NATsess DevTyp="AFTR" DevID="bgw211.example.net"
SiteID="A2E0:62" PostS4="198.51.100.127"
Proto="6" PreSPt="49156" PostSPt="6083" TrigR="I"]
```

Character count: about 205.

The next example is perhaps more typical in size. Assume an enterprise NAT44 generating its own logs. The optional parameters are omitted. This is a session deletion event.

```
<86>1 2013-05-07T15:27:49.603-04:00 cerberus.example.com
NAT 175 SessDel [NATsess SiteID="192.0.2.5" PostS4="198.51.100.14"
Proto="6" PreSPt="51387" PostSPt="17865"]
```

The character count: about 165.

5.3.2. Address Binding Event

As shown in Table 1, the NAT address binding event is indicated by MSG-ID set to "AddrBind". The associated SD-ELEMENT is tagged by SD-ID "NATBind". The contents of the NATBind SD-ELEMENT are shown in Table 4. The requirements for these contents are derived from the description in [Section 3.3](#).

PARAM-NAME	Description	Requirement
DevTyp	Section 5.2.3	OPTIONAL
DevID	Section 5.2.2	OPTIONAL
SiteID	Section 5.2.11	MANDATORY

PostS4	Section 5.2.4	MANDATORY	
+-----+	+-----+	+-----+	+-----+

Table 4: Contents Of the SD-ELEMENT Section For Logging the Address Binding Event

As an example, consider a DS-Lite AFTR [[RFC6333](#)] incorporating a PCP server, where PCP is used to obtain an external address binding and a port range. See [Section 11 of \[RFC6887\]](#) for the address binding. (The port allocation is shown in the next section's example.) As in the session creation example, the first /16 prefix and the final 64 bits are omitted from the encapsulating IPv6 address which is used as the subscriber site identifier.

```
<86>1 2013-05-07T15:27:49.603Z yourd137mzmhow.example.net
NAT 68 AddrBind [NATBind SiteID="5A27:876E" PostS4="198.51.100.1"]
```

Character count: about 125.

5.3.3. Port Allocation Change

As indicated in Table 1, the port block allocation change event is indicated by MSG-ID set to "PtAlloc". The associated SD-ELEMENT is tagged by SD-ID "NATPBlk". The contents of the NATPBlk SD-ELEMENT are shown in Table 5. The requirements for these contents are derived from the description in [Section 3.4](#).

+-----+	+-----+	+-----+	+-----+
PARAM-NAME	Description	Requirement	
+-----+	+-----+	+-----+	+-----+
DevTyp	Section 5.2.3	OPTIONAL	
DevID	Section 5.2.2	OPTIONAL	
SiteID	Section 5.2.11	MANDATORY	
PostS4	Section 5.2.4	MANDATORY	
Ptrg	Section 5.2.10	MANDATORY	
+-----+	+-----+	+-----+	+-----+

Table 5: Contents Of the SD-ELEMENT Section For Logging the Port Allocation Change Event

As in the example in the previous section example, consider a DS-Lite AFTR [[RFC6333](#)] incorporating a PCP server, where PCP is used to obtain an external address binding and a port range. See [\[I-D.pcp-port-set\]](#) for the port set part of this operation.

Strictly for purposes of illustration, assume that the subscriber is allocated two ranges of 64 consecutive values each, with the first beginning at 2048 and the second at 4096.


```
<86>1 2013-05-07T15:27:49.751Z yourd137mzmhow.example.net
NAT 68 PtAlloc [NATPBlk SiteID="5A27:876E" PostS4="198.51.100.1"
PtRg="2048-2111" PtRg="4096-4159"]
```

Character count: about 155.

5.3.4. Address Exhaustion Event

As indicated in Table 1, the address exhaustion event is indicated by MSG-ID set to "AddrEx". The associated SD-ELEMENT is tagged by SD-ID "NATAddrEx". The contents of the NATAddrEx SD-ELEMENT are shown in Table 6. The requirements for these contents are derived from the description in [event deleted].

PARAM-NAME	Description	Requirement
DevTyp	Section 5.2.3	OPTIONAL
DevID	Section 5.2.2	OPTIONAL
APoolId	Section 5.2.1	MANDATORY

Table 6: Contents Of the SD-ELEMENT Section For Logging the Address Exhaustion Event

The example shows this event being reported by a DS-Lite AFTR. Note the critical priority indication at the beginning of the log. As with the session example, we assume off-board log generation.

```
<82>1 2013-05-07T22:14:15.03Z record.example.net NAT 5063
AddrEx [NATAddrEx DevID="bgw211.example.net" APoolId="2"]
```

Character count: about 120.

5.3.5. NAT Port Exhaustion

As indicated in Table 1, the port exhaustion event is indicated by MSG-ID set to "PortEx". The associated SD-ELEMENT is tagged by SD-ID "NATPEX". The contents of the NATPEX SD-ELEMENT are shown in Table 7. The requirements for these contents are derived from the description in [event deleted].

PARAM-NAME	Description	Requirement
DevTyp	Section 5.2.3	OPTIONAL
DevID	Section 5.2.2	OPTIONAL
PostS4	Section 5.2.4	MANDATORY

Proto	Section 5.2.7	MANDATORY	
+-----+	+-----+	+-----+	+-----+

Table 7: Contents Of the SD-ELEMENT Section For Logging the Port Exhaustion Event

The example is straightforward. Note the warning priority indication at the beginning of the log.

```
<84>1 2013-05-07T22:14:15.03Z cerberus.example.com NAT 5063
PortEx [NATPEX PostS4="198.51.100.1" Proto="6"]
```

Character count: about 110.

5.3.6. Quota Exceeded

As indicated in Table 1, the quota exceeded event is indicated by MSG-ID set to "Quota". The associated SD-ELEMENT is tagged by SD-ID "NATQEx". The contents of the NATQEx SD-ELEMENT are shown in Table 8. The requirements for these contents are derived from the description in [Section 3.13](#).

PARAM-NAME	Description	Requirement	
+-----+	+-----+	+-----+	+-----+
DevTyp	Section 5.2.3	OPTIONAL	
DevID	Section 5.2.2	OPTIONAL	
SScop	Section 5.2.12	MANDATORY	
PScop	Section 5.2.8	MANDATORY	
SiteID	Section 5.2.11	OPTIONAL	
VLANid	Section 5.2.14	OPTIONAL	
VRFid	Section 5.2.15	OPTIONAL	
+-----+	+-----+	+-----+	+-----+

Table 8: Contents Of the SD-ELEMENT Section For Logging the Quota Exceeded Event

Example 1: limit on TCP sessions for a specific user site reached at an AFTR with off-board log generation.

```
<85>1 2013-05-07T22:14:15.03Z record.example.net NAT 5063
Quota [NATQEx DevID="bgw211.example.net" SScop="S" PScop="6"
SiteID="A2E0:62"]
```

Character count: about 135.

Example 2: global limit on number of sessions for all subscribers served by the same VLAN.


```
<85>1 2013-05-07T15:27:49.603-04:00 cerberus.example.com
NAT 175 Quota [NATQEx SScop="M" PScop="*" VLANid="1246"]
```

Character count: about 115.

Example 3: limit on total number of sessions for TCP.

```
<85>1 2013-05-07T15:27:49.603-04:00 cerberus.example.com
NAT 175 Quota [NATQEx SScop="*" PScop="6"]
```

Character count: about 100.

6. IANA Considerations

This document requests IANA to make the following assignments to the SYSLOG Structured Data ID Values registry. RFCxxxx refers to the present document when approved.

Structured Data ID	Structured Data Parameter	Required or Optional	Reference
NATsess		OPTIONAL	RFCxxxx
	DevTyp	OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	SiteID	MANDATORY	RFCxxxx
	PostS4	MANDATORY	RFCxxxx
	Proto	MANDATORY	RFCxxxx
	PreSPt	MANDATORY	RFCxxxx
	PostSPt	MANDATORY	RFCxxxx
	TrigR	OPTIONAL	RFCxxxx
----	----	----	----
NATBind		OPTIONAL	RFCxxxx
	DevTyp	OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	SiteID	MANDATORY	RFCxxxx
	PostS4	MANDATORY	RFCxxxx
----	----	----	----
NATPBlk		OPTIONAL	RFCxxxx
	DevTyp	OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	SiteID	MANDATORY	RFCxxxx
	PostS4	MANDATORY	RFCxxxx
	PtRg	MANDATORY	RFCxxxx
----	----	----	----
NATAddrEx		OPTIONAL	RFCxxxx
	DevTyp	OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx

	APoolId	MANDATORY	RFCxxxx
----	----	----	----
NATPEX		OPTIONAL	RFCxxxx
	DevTyp	OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	PostS4	MANDATORY	RFCxxxx
	Proto	MANDATORY	RFCxxxx
----	----	----	----
NATQEX		OPTIONAL	RFCxxxx
	DevTyp	OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	SScop	MANDATORY	RFCxxxx
	PScop	MANDATORY	RFCxxxx
	SiteID	OPTIONAL	RFCxxxx
	VLANid	OPTIONAL	RFCxxxx
	VRfid	OPTIONAL	RFCxxxx
----	----	----	----
NATInvP		OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	SiteID	MANDATORY	RFCxxxx
	PSID	MANDATORY	RFCxxxx
+-----+-----+-----+-----+			

Table 9: NAT-Related STRUCTURED-DATA Registrations

7. Management Considerations

To come.

8. Security Considerations

When logs are being recorded for regulatory reasons, preservation of their integrity and authentication of their origin is essential. To achieve this result, it is RECOMMENDED that the operator deploy [\[RFC5848\]](#).

Access to the logs defined here while the reported assignments are in force could improve an attacker's chance of hijacking a session through port-guessing. Even after an assignment has expired, the information in the logs SHOULD be treated as confidential, since, if revealed, it could help an attacker trace sessions back to a particular subscriber or subscriber location. It is therefore RECOMMENDED that these logs be transported securely, using [\[RFC5425\]](#), for example, and that they be stored securely at the collector.

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.
- [RFC2685] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", [RFC 2685](#), September 1999.
- [RFC5424] Gerhards, R., "The Syslog Protocol", [RFC 5424](#), March 2009.
- [RFC5425] Miao, F., Ma, Y., and J. Salowey, "Transport Layer Security (TLS) Transport Mapping for Syslog", [RFC 5425](#), March 2009.
- [RFC5848] Kelsey, J., Callas, J., and A. Clemm, "Signed Syslog Messages", [RFC 5848](#), May 2010.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", [RFC 5952](#), August 2010.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", [RFC 6145](#), April 2011.
- [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

- [I-D.Behave-NAT-MIB]
Perreault, S., Tsou, T., and S. Sivakumar, "Additional Managed Objects for Network Address Translators (NAT) (Work in progress)", July 2013.
- [I-D.behave-ipfix-nat-logging]
Sivakumar, S. and R. Penno, "IPFIX Information Elements for logging NAT Events (Work in progress)", March 2013.
- [I-D.pcp-port-set]
Sun, Q., Boucadair, M., Sivakumar, S., Zhou, C., Tsou, T., and S. Perreault, "Port Control Protocol (PCP) Extension for Port Set Allocation (Work in progress)", March 2013.
- [I-D.softwire-lw4over6]

Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture (Work in progress)", April 2013.

[I-D.softwire-map]

Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., and T. Murakami, "Mapping of Address and Port with Encapsulation (MAP) (Work in progress)", March 2013.

[I-D.softwire-public-4over6]

Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public IPv4 over IPv6 Access Network (Work in progress)", February 2013.

[I-D.softwire-unified-cpe]

Boucadair, M. and I. Farrer, "Unified IPv4-in-IPv6 Softwire CPE (Work in progress)", March 2013.

[I-D.tsou-behave-natx4-log-reduction]

Tsou, T., Li, W., and T. Taylor, "Port Management To Reduce Logging In Large-Scale NATs (Work in progress)", May 2013.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", [BCP 127](#), [RFC 4787](#), January 2007.

[RFC5101] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", [RFC 5101](#), January 2008.

[RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", [BCP 142](#), [RFC 5382](#), October 2008.

[RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", [RFC 5969](#), August 2010.

[RFC6264] Jiang, S., Guo, D., and B. Carpenter, "An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition", [RFC 6264](#), June 2011.

- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [RFC 6333](#), August 2011.
- [RFC6674] Brockners, F., Gundavelli, S., Speicher, S., and D. Ward, "Gateway-Initiated Dual-Stack Lite Deployment", [RFC 6674](#), July 2012.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [RFC 6887](#), April 2013.
- [RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", [BCP 127](#), [RFC 6888](#), April 2013.

Authors' Addresses

Zhonghua Chen
China Telecom
P.R. China

Email: 18918588897@189.cn

Cathy Zhou
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: cathy.zhou@huawei.com

Tina Tsou
Huawei Technologies (USA)
2330 Central Expressway
Santa Clara, CA 95050
USA

Phone: +1 408 330 4424

Email: tina.tsou.zouting@huawei.com

T. Taylor (editor)
Huawei Technologies
Ottawa
Canada

Email: tom.taylor.stds@gmail.com