Behave Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 29, 2014

Z. Chen
China Telecom
C. Zhou
T. Tsou
T. Taylor, Ed.
Huawei Technologies
January 25, 2014

Syslog Format for NAT Logging
draft-ietf-behave-syslog-nat-logging-06

Abstract

NAT devices are required to log events like creation and deletion of translations and information about the resources the NAT is managing. The logs are required to identify an attacker or a host that was used to launch malicious attacks, and for various other purposes of accounting and management. Since there is no standard way of logging this information, different NAT devices behave differently. The lack of a consistent way makes it difficult to write the collector applications that would receive this data and process it to present useful information.

This document describes the information that is required to be logged by the NAT devices. 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 7011). 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.

Table of Contents

1. Introduction 4
1.1. Terminology 5
2. Deployment Considerations 6
2.1. Static and Dynamic NATs 6
2.2. Realms and Address Pools 7
2.2.1. Address Pools 7
2.3. NAT Logging Requirements For Different Transition Methods 8
2.4. Subscriber Identification 9
2.5. The Port Control Protocol (PCP) 10
2.6. Logging At the Customer Edge 10
3. NAT-Related Events and Parameters 10
3.1. Events Relating To Allocation Of Resources To Hosts . . . 10
3.1.1. NAT Address Mapping Creation and Deletion 11
3.1.2. NAT Address and Port Mapping Creation and Deletion . 12
3.1.3. NAT Session Creation and Deletion 14
3.1.3.1. Destination Logging 17
3.1.4. Port Range Allocation and Deallocation 17
3.2. Threshold Events 19
3.2.1. Address Pool High- and Low-Water-Mark Threshold Events 19
3.2.2. Global Address Mapping High-Water-Mark Threshold Event 20
3.2.3. Global Address and Port Mapping High-Water-Mark Threshold Event 21
3.2.4. Subscriber-Specific Address and Port Mapping Threshold Event 22
3.3. Limit-Related Events 22
3.3.1. Global Address Mapping Limit Exceeded 22
3.3.2. Global Address and Port Mapping Limit Exceeded . . . 23
3.3.3. Global Limit On Number of Active Hosts Exceeded . . . 24
3.3.4. Subscriber-Specific Limit On Number of Address and

Port Mappings Exceeded 25
3.3.5. Global Limit On Number Of Fragments Pending Reassembly Exceeded 26
4. SYSLOG Applicability 27
5. SYSLOG Record Format For NAT Logging 27
5.1. SYSLOG HEADER Fields 28
5.2. Parameter Encodings 29
5.2.1. General Encoding Rules 32
5.2.2. Special Cases 32
5.2.3. Relationship To Objects In the NAT MIB 33
5.3. Encoding Of Complete Log Report For Each Event Type . . . 35
5.3.1. Encoding of Events Relating To Allocation Of Resources To Hosts 35
5.3.1.1. NAT Address Mapping Creation and Deletion 36
5.3.1.2. NAT Address and Port Mapping Creation and Deletion 37
5.3.1.3. NAT Session Creation and Deletion 39
5.3.1.4. Port Range Allocation and Deallocation 41
5.3.2. Encoding of Threshold Events 43
5.3.2.1. NAT Address Pool High- and Low-Water-Mark Threshold Events 43
5.3.2.2. Global Address Mapping High-Water-Mark Threshold Exceeded 44
5.3.2.3. Global Address and Port Mapping High-Water-Mark Threshold Event 45
5.3.2.4. Subscriber-Specific Address and Port Mapping High-Water-Mark Threshold Event 45
5.3.3. Encoding of Limit Events 46
5.3.3.1. Global Address Mapping Limit Exceeded 46
5.3.3.2. Global Address and Port Mapping Limit Exceeded . 47
5.3.3.3. Global Limit On Number of Active Hosts Exceeded . 48
5.3.3.4. Subscriber-Specific Limit On Number of Address and Port Mappings Exceeded 49
5.3.3.5. Pending Fragment Limit Exceeded 50
6. Management Considerations 51
6.1. General Requirements For Control Of Logging 51
6.1.1. Configuration of PRI Value 51
6.1.2. Ability For Each Collector To Detect Lost Event Reports 52
6.1.3. Ability To Rate Limit Or Disable Event Reports . . . 52
6.2. Setting Limits and Thresholds 53
6.3. Other Management Requirements 54
7. Security Considerations 55
8. IANA Considerations 55
9. References 58
9.1. Normative References 58
9.2. Informative References 60
Authors' Addresses 61

1. Introduction

This document deals with logging of NAT activity in two categories: NAT translations and NAT resource usage.

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 that 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 [RFC7011] respectively. The same content is proposed to be logged by both documents.

In addition to records of subscriber activity, some operators use logs to indicate when utilization of critical resources is approaching or has reached limits set by the operator or implementation. This document and the IPFIX document therefore provide logs in two categories: thresholds exceeded and limits exceeded. Operators have the alternative to receive the threshold limits as SNMP notifications (see the NAT MIB [I-D.behave-NAT-MIB]).

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 detailed description of the events that need logging and the parameters that may be required in the logs. Section 3.1 describes events related to subscriber activity, Section 3.2 covers threshold events, and Section 3.3 covers events where hard limits have been reached.

The use of SYSLOG [RFC5424] has advantages and disadvantages compared with the use of IPFIX [RFC7011]. 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. Section 5.1 describes the SYSLOG header format for each report, Section 5.2 lists and describes the encoding of parameters that can appear in the logs, and Section 5.3 specifies the encoding of the body of each event report. 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 makes frequent reference to the NAT MIB. That reference is to the document [I-D.behave-NAT-MIB].

This document makes frequent reference to NAT behaviours defined in [RFC4784]. In particular it refers to

- the recommended pooling behaviour "pooled" and its contrary pooling behaviour "arbitrary"; and

- the recommended mapping behaviour "endpoint-independent" and its contrary mapping behaviour "endpoint-dependent".

This document uses the term "address mapping" to denote an association between an internal IP address and an IP address in a selected external realm. See Section 2.2 for a further discussion of this process.

> The natMapIntAddrTable in the NAT MIB provides details on all currently active address mappings. Note that this table is applicable only when NAT pooling behaviour is "paired".

This document uses the [RFC4787] term "address and port mapping" to denote a three-tuple association between an internal IP address and port and an IP address and port in a selected external realm, or between an internal <IP address, ICMP identifier> pair and an <IP address, ICMP identifier> pair in the selected realm. For implementations which maintain a Binding Information Base (BIB) (as described in Section 2 of [RFC6146], for example), the content of a BIB entry is an address and port mapping.

> The natMappingTable in the NAT MIB provides details on all currently active address and port mappings.

This document uses the term "session" as it is defined in [RFC2663], Section 2.3. From the point of view of this document, session creation involves the combination of a source address and port mapping with a mapping between internal and external destination address and port to create a full five-tuple mapping.

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.

To match the terminology used by the NAT MIB, this document uses the term "subscriber" to denote any device being served by the NAT, whether individual host or customer edge router. That is, despite the carrier-oriented terminology, the intended scope of applicability of this document is both to NATs in the carrier network and managed NATs in the customer network.

Finally, with two exceptions, when the terms "source" or "destination" are used below, they denote the source and destination of packets that are flowing from the internal to the external realm, regardless of the direction of session establishment or the direction of flow of an individual packet. The exceptions relate to the global address and port mapping limit event and the pending fragment limit event, when the actual source and destination addresses in the header of the packet that hit the limit are reported.

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.

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 address and port 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 identifies an IP numbering space. A NAT session always maps between an internal and an external realm. In simple NAT configurations, it may be possible to identify a default internal realm and/or a default external realm for all sessions. In more complex NAT configurations a given realm may be an internal realm for some sessions and an external realm for others. Realms without subscriber sites are always external.

Address pools are associated with specific realms in their external role.

It is necessary to define multiple realms when the NAT supports overlapping IP numbering spaces. In such a case, the NAT must determine the source realm and subscriber using additional information associated with the incoming packet. See further discussion in Section 2.4.

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 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 mappings depends on the pooling behaviour. With a pooling behaviour of "arbitrary", 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", 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 3.2.1 that allow monitoring of port utilization at the pool level. Section 6.2 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 [RFC7040], 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 will not be unique. As a consequence, identification of the realm and subscriber from which the packet was sent needs to include an additional 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], two identifiers are needed: an identifier of the softwire from the gateway to the NAT, and an identifier associated with the incoming tunnel to the gateway.

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 edge. 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.

2.4. Subscriber Identification

The ability to identify the particular subscriber involved in an event is required for the events defined in Section 3.1, and desirable for technician follow-up for those defined in Section 3.2.4 and Section 3.3.

As mentioned above, in some NAT configurations the source address is insufficient to identify an individual subscriber because of overlapping address space, and additional information is required. For example, if the NAT supports DS-Lite [RFC 6333], the source address of incoming packets from DS-Lite subscribers will always be in the range 192.0.0/29. The additional information required in this case is the IPv6 address of the encapsulating header.

The natSubscribersTable in the NAT MIB contains the additional information needed, if any, to identify each subscriber. Thus it is sufficient to include the index to this table in the event report to provide the needed identification. However, this implies that for full interpretation of the event report, the configuration information stored in the natSubscribersTable must be stored (along with AAA information relating the additional identifiers to the subscriber profiles, which must be stored in any event). To relieve the operator of the need to store the configuration data (given that the logs may be needed months or years after they were recorded), the reports specified in Section 3.1 include the additional identifying information that is found in the natSubscribersTable.

This document standardizes the presentation of the following possible additional classifying information within NAT-related log reports:

- interface index [RFC2863];
- VLAN index [RFC4363];
- VPN identifier [RFC4265];
- DS-Lite encapsulating IPv6 address [RFC6333].

Which of these is actually used in a given NAT depends on implementation and deployment.

Gateway-Initiated DS-Lite [RFC6674] identifiers could also be specified, but it seems premature to do so because it is not clear which of the variety of possibilities presented in Section 6 and Appendix A.2 of [RFC6674] are actually being deployed.

2.5. 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.

2.6. 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. Section 4 of [RFC6888] also provides a brief statement of logging requirements for carrier grade NATs.

In SYSLOG, 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.

Parameters marked CONDITIONAL are REQUIRED under some circumstances but not others. Details are provided for each event.

3.1. Events Relating To Allocation Of Resources To Hosts

Setting up a NAT session proceeds in a series of logical steps: creation of an address mapping, creation of an address and port mapping, and finally, creation of the session.

The reports corresponding to these three steps are defined in Section 3.1.1, Section 3.1.2, and Section 3.1.3 respectively. Which of these reports is enabled depends on the NAT implementation and operator preferences, subject to the considerations of the next paragraph.

If the NAT implements the recommended pooling behaviour of "paired", address mapping creation is an event distinct in general from the creation of a subsequent address and port mapping based on that address mapping. However, if the pooling behaviour is "arbitrary" [RFC4787], the two events occur simultaneously and there is no point in reporting both. Similarly, if the NAT implements the recommended mapping behaviour of "endpoint-independent mapping", the two events of address and port mapping creation and session creation based on that mapping are distinct and may meaningfully be reported separately. However, if the mapping behaviour is "endpoint-dependent", the two events occur simultaneously and it is only meaningful to report session creation.

The fourth report type in this section describes the bulk allocation of ports to an address mapping, which the NAT may implement if the pooling behaviour is "paired" [RFC4787]. It, along with the other reports, is needed to provide complete accountability for resources allocated to the subscriber.

3.1.1. NAT Address Mapping Creation and Deletion

Two specific events are provided:

- NAT address mapping creation;
- NAT address mapping deletion.

Implementations MUST NOT report these events unless pooling behaviour is "paired".

Address mapping is discussed in detail in Section 2.2.

One address mapping creation event is associated with potentially many succeeding address and port mapping creation events, as individual port values are mapped for specific protocols. Similarly, an address mapping deletion event may be associated with potentially many address and port mapping 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:

- NAT instance identifier (CONDITIONAL);
- Source subscriber index (MANDATORY);
- Additional source subscriber classifier value as recognized at the ingress to the internal realm (CONDITIONAL);
- Internal realm (CONDITIONAL);
- Internal address type (MANDATORY);
- Internal source IP address (MANDATORY);
- External realm (CONDITIONAL);
- External address type (MANDATORY);
- External source IP address (MANDATORY);
- 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.

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.
- Additional source subscriber classifier value REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.
- Internal or external realm REQUIRED if not the default internal or external realm, respectively, else MAY appear.

3.1.2. NAT Address and Port Mapping Creation and Deletion

The address and port mapping creation or deletion event reports the addition or deletion of an address and port mapping as defined in Section 1.1. If the implementation maintains a Binding Information Base (BIB), this is equivalent to the creation or deletion of a BIB entry. Implementations MUST support the generation of the address and port mapping creation/deletion event reports if they implement

the recommended mapping behaviour "endpoint-independent". They MAY support reporting of these events in the contrary case.

The address and port mapping creation/deletion event report provides the same information as the session creation/deletion event, except for the destination-related fields and (in general) timestamp values in the latter. With "endpoint-independent" mapping behaviour, one address and port mapping creation event is associated with potentially many succeeding session creation events. Similarly, an address and port mapping 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 address and port mapping deletion.

Operators should disable the reporting of address and port mapping creation and deletion events when destination logging is enabled, because of the redundancy between the address and port mapping and session event reports. However, if destination logging is disabled and the NAT uses the recommended "endpoint-independent" mapping behaviour, it is the session events that are redundant and should be disabled.

The following specific events are defined:

- NAT address and port mapping creation

- NAT address and port mapping deletion

These take the same parameters for all types of NAT. The internal realm, subscriber-identifying information, internal source IP address, external realm, and external source IP address capture the underlying address mapping. The port values and protocol are unique to the address and port mapping.

The parameters for the address and port mapping creation/deletion event are:

- NAT instance identifier (CONDITIONAL);

- Source subscriber index (MANDATORY);

- Additional source subscriber classifier value as recognized at the ingress to the internal realm (CONDITIONAL);

- Internal realm (CONDITIONAL);

- Internal address type (MANDATORY);

- Internal source IP address (MANDATORY);
- Internal source port or ICMP identifier (MANDATORY);
- External realm (CONDITIONAL);
- External address type (MANDATORY);
- External source IP address (MANDATORY);
- External source port or ICMP identifier (MANDATORY);
- Protocol identifier (MANDATORY);
- Trigger for address and port 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 (applicable only if pooling behaviour is "paired" [RFC4787]).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.
- Additional source subscriber classifier value REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.
- Internal or external realm REQUIRED if not the default internal or external realm, respectively, else MAY appear.

3.1.3. NAT Session Creation and Deletion

A NAT session creation or deletion event is logged when a address and port mapping 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 address and port mapping.

Implementations MUST provide a means for the operator to specify whether destination information is to be included in the reports of these events (see discussion below).

The following specific events are defined:

- NAT session creation
- NAT session deletion

These take the same parameters for all types of NAT. Parameters "internal realm" through "protocol identifier" capture the underlying address and port mapping. Subsequent parameters capture the destination address and destination subscriber identity (if applicable).

The parameters for the session creation/deletion event are:

- NAT instance identifier (CONDITIONAL);
- Internal source subscriber index, equal to the natSubscriberIndex value in the natSubscribersTable in the NAT MIB (MANDATORY);
- Additional internal subscriber classifier value (CONDITIONAL);
- Internal realm (CONDITIONAL);
- Internal address type (MANDATORY);
- Internal source IP address (MANDATORY);
- Internal source port or ICMP identifier (MANDATORY);
- External realm (CONDITIONAL);
- External address type (MANDATORY);
- External source IP address (MANDATORY);
- External source port or ICMP identifier (MANDATORY);
- Protocol identifier (MANDATORY);
- Internal destination IP address (CONDITIONAL);
- Internal destination port or ICMP identifier (CONDITIONAL);
- Destination subscriber index (CONDITIONAL);

- Additional destination subscriber classifier value as recognized at the ingress to the external realm (CONDITIONAL);
- External destination IP address (CONDITIONAL);
- External destination port or ICMP identifier (CONDITIONAL);
- 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 address and port mapping (applicable only if the NAT mapping behaviour is "endpoint-independent".

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.
- Additional source subscriber classifier value REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.
- Internal or external realm REQUIRED if not the default internal or external realm, respectively, else MAY appear.
- Internal destination address and port REQUIRED if destination logging is enabled and these need to be remapped to external destination address and port. Otherwise, if destination logging is disabled, they MUST NOT appear, and if destination logging is enabled, they SHOULD NOT appear because of redundancy.
- External destination subscriber index REQUIRED if destination logging is enabled and the destination is a subscriber served by the NAT, else MUST NOT appear.
- Additional external subscriber classifier value REQUIRED if destination logging is enabled and the destination is a subscriber served by the NAT and the external destination address is not enough to identify the external destination subscriber unambiguously, else MUST NOT appear.

- External destination address and port REQUIRED if destination logging is enabled, else MUST NOT appear.

3.1.3.1. Destination Logging

The logging of destination address and port is 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 destination information for sessions matching the specified criteria, but MUST NOT report these events for other sessions.

3.1.4. Port Range 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:

- Port range allocation;
- Port range deallocation.

The parameters for these events are:

- NAT instance identifier (CONDITIONAL);
- Source subscriber index (MANDATORY);
- Additional source subscriber classifier value as recognized at the ingress to the internal realm (CONDITIONAL);

- Internal realm (CONDITIONAL);
- Internal address type (MANDATORY);
- Internal source IP address (MANDATORY);
- External realm (CONDITIONAL);
- External address type (MANDATORY);
- External source IP address (MANDATORY);
- Lowest port number of the range being allocated or deallocated (MANDATORY).
- Highest port number of the range being allocated or deallocated (MANDATORY).
- Trigger for port range allocation or deallocation (OPTIONAL):
  - outgoing packet received;
  - incoming packet received;
  - administrative action (e.g., via the Port Control Protocol [RFC6887]); or
  - autonomous action of the NAT.

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.
- Additional source subscriber classifier value REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.
- Internal or external realm REQUIRED if not the default internal or external realm, respectively, else MAY appear.

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

3.2. Threshold Events

The events of this section are based on thresholds set by the operator within the NAT MIB. Cross-references to the associated MIB objects are provided for each event. With the exception of the address pool low-water-mark event, the threshold events provide early warning of potential dropped packets due to resource exhaustion or administrator-imposed limits.

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

Two specific events provide reports on address pool utilization:

- High-water-mark threshold reached or exceeded;

- 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 rather than logging these events.

Address pools are discussed in Section 2.2.1. The natPoolTable object in the NAT MIB provides access to parameters describing the utilization level of address-port combinations within a given pool. Since a new 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 requested 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 6.2 discusses factors affecting the choice of the threshold values.

The high-water-mark threshold event provides a warning that the address-port combinations offered by the pool are nearing exhaustion. Upon exhaustion, subscribers may be unable to establish new connections because no address has enough free port values left to be allocated to an address mapping ("address exhaustion"). This applies to the case of "paired" pooling behaviour, where typically an address

will not be allocated unless it has a sufficient number of free ports. Alternatively, new connections cannot be established simply because no address in the pool has a free port number for the required protocol ("port exhaustion").

Packets triggering failed attempts to establish new connections due to address exhaustion are included in the following NAT MIB dropped packet counters:

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

Packets triggering failed attempts to establish new connections due to port exhaustion are counted in the following NAT MIB dropped packet counters:

- globally, natOutOfPortErrors in the natCounters table;
- per protocol, natProtocolOutOfPortErrors in natProtocolTable;
- per subscriber, natSubscriberOutOfPortErrors in natSubscribersTable.

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

- NAT instance identifier (CONDITIONAL);
- Pool identifier (MANDATORY);
- The threshold value set by the administrator (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

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

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

- Address mapping high-water-mark threshold exceeded.

Implementations MUST NOT generate this event report unless the pooling behaviour is "paired". Depending on deployment, operators can choose instead to use the SNMP notification natNotifAddrMappings defined in the 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 is greater than 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:

- NAT instance identifier (CONDITIONAL);
- Current number of active address mappings (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

3.2.3. Global Address and Port Mapping High-Water-Mark Threshold Event

One specific event allows monitoring of the total number of active address and port mappings. Where the NAT implements a BIB, this is equivalent to the total number of BIB entries.

- address and port mapping high-water-mark threshold exceeded.

Depending on deployment, operators can choose instead to use the SNMP notification natNotifMappings defined in the NAT MIB.

The NAT MIB displays cumulative counts of address and port mappings created and removed in the natCounters table. When the difference between these two counters is greater than 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:

- NAT instance identifier (CONDITIONAL);
- Current number of active address and port mappings (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

3.2.4. Subscriber-Specific Address and Port Mapping Threshold Event

An event is provided to allow monitoring of the total number of active mappings per subscriber:

- Subscriber-specific mapping high-water-mark threshold exceeded.

Depending on deployment, operators can choose instead to use the SNMP notification natNotifSubscriberMappings defined in the NAT MIB.

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

The specific parameters provided by this event report are:

- NAT instance identifier (CONDITIONAL);
- Index of the affected subscriber (MANDATORY).
- Current number of active mappings for this subscriber (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

3.3. Limit-Related Events

The events of this section are generated when hard limits set by the operator are exceeded. The consequence for service will be dropped packets. As with the threshold events, the description of each report includes cross-references to the associated MIB objects.

3.3.1. Global Address Mapping Limit Exceeded

The global address mapping limit exceeded event is reported when a new address mapping is requested 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:

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

Implementations MUST NOT generate this event report unless the pooling behaviour is "paired". Depending on deployment, operators can choose instead to use the SNMP notification natNotifAddrMappings defined in the NAT MIB.

The parameters for this event are:

- NAT instance identifier (CONDITIONAL);
- Index of the affected subscriber (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

The subscriber index is provided to allow the operator to correlate the event with any subscriber complaints or possible abuse.

3.3.2. Global Address and Port Mapping Limit Exceeded

The global address and port mapping limit exceeded event is reported when a new address and port mapping is requested but the total number of address and port 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:

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

The parameters for this event are:

- NAT instance identifier (CONDITIONAL);
- Index of the internal subscriber (CONDITIONAL);

- Index of the external subscriber (CONDITIONAL);
- Source realm of the triggering packet (MANDATORY);
- Incoming packet header IP address type (CONDITIONAL);
- Incoming packet source IP address (CONDITIONAL).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.
- The index of the internal subscriber is REQUIRED if the mapping was triggered by a packet outgoing from the internal to the external realm, else MUST NOT appear.
- The index of the external subscriber is REQUIRED if the mapping was triggered by a packet incoming from a subscriber served by the NAT and located in the external realm (i.e., using an address mapping created previously by the internal subscriber), else MUST NOT appear.
- The address type and source IP address from the initiating packet are REQUIRED if the mapping was triggered by a packet incoming from a purely external realm (i.e., using an address mapping created previously by the internal subscriber), else MAY appear.

The subscriber index or packet source address is provided to allow the operator to correlate the event with any subscriber complaints or possible abuse.

3.3.3. Global Limit On Number of Active Hosts Exceeded

The global limit on number of active hosts exceeded event is reported when an address mapping is requested (at least at the logical level) for a host with no previous active mappings, but the total number of active hosts 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:

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

The parameters for this event are:

- NAT instance identifier (CONDITIONAL);
- Index of the affected subscriber (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

The subscriber index is provided to allow the operator to correlate the event with any subscriber complaints.

3.3.4. Subscriber-Specific Limit On Number of Address and Port Mappings Exceeded

The subscriber-specific limit on number of address and port mappings exceeded event is reported when a new mapping is requested, but the total number of active mappings 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:

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

The parameters for this event are:

- NAT instance identifier (CONDITIONAL);
- Index of the affected subscriber (MANDATORY).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

The subscriber index is provided to allow the operator to take administrative action or to correlate the event with any subscriber complaints or possible abuse.

3.3.5. 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:

- globally, natResourceErrors in the natCounters table;
- per protocol, natProtocolResourceErrors in natProtocolTable;
- 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 of the packet is a subscriber served by the NAT and the subscriber index can be determined, it MUST also be included.

- NAT instance identifier (CONDITIONAL);
- Source realm of the packet (MANDATORY);
- Packet header IP address type (MANDATORY);
- Packet source IP address (MANDATORY);
- Packet destination IP address (MANDATORY);
- Source subscriber index (CONDITIONAL).

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.
- Source subscriber index REQUIRED if the source of the packet is a subscriber served by the NAT and can be determined, else MUST NOT appear.

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 other header 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 example in Section 5.3.1.3, if destination logging is enabled, typical record sizes for session event logs are in the order of 300 bytes, so throughput capacity should be higher than in the call detail case for the same amount of computing power. However, note that bursts of session deletion events may occur as a result of deletion of the underlying mapping or address mapping.

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 Severity value of 6 (Informational) is suggested for the events relating to creation and deletion of sessions, mappings, address mappings, and port allocation, combined with a suitable Facility value in the range 16-23 (local use) to ensure routing to a secure collector. The Facility value(s) for the threshold and limit events will presumably be chosen to route them to maintenance for immediate action and/or to provisioning for less urgent consideration. The suggested value of Severity by event type is shown in Table 1, but in practice has a clear dependency on the context within which the NAT is operating.

The TIMESTAMP field SHOULD be expressed with sufficient precision to distinguish non-simultaneous event occurrences, subject to the accuracy of the local clock. This specification does not assume the ability to correlate the events reported by the subject device with events recorded by other devices, although that may be required for other reasons. Hence from the point of view of this specification only relative rather than absolute accuracy is of interest.

The HOSTNAME header field MUST identify the NAT device. 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. To simplify log collection procedures, the APP-NAME value "NAT" MUST be used for the event reports specified in Section 5.3.1, the APP-NAME value "NATTHR" MUST be used for the event types defined in Section 5.3.2, and the APP-NAME value "NATLIM" MUST be used for the event types defined in Section 5.3.3.

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	APP-NAME	MSGID	Severity	SD-ID
NAT address mapping creation	NAT	AMADD	6 info	namap
NAT address mapping deletion	NAT	AMDEL	6 info	namap
NAT address and port mapping creation	NAT	APMADD	6 info	napmap
NAT address and port mapping deletion	NAT	APMDEL	6 info	napmap
NAT session creation	NAT	SADD	6 info	nsess
NAT session deletion	NAT	SDEL	6 info	nsess
Port range allocation	NAT	PTADD	6 info	nprng
Port range deallocation	NAT	PTDEL	6 info	nprng
Address pool high threshold	NATTHR	POOLHT	4 warning	npool
Address pool low threshold	NATTHR	POOLLT	6 info	npool
Global address mapping high threshold	NATTHR	GAMHT	4 warning	ngamht
Global address and port mapping high threshold	NATTHR	GAPMHT	4 warning	ngapmht
Subscriber-specific mapping high threshold	NATTHR	SAPMHT	5 notice	nsapmht
Global address mapping limit	NATLIM	GAMLIM	3 error	ngaml
Global address and port mapping limit	NATLIM	GAPMLIM	3 error	ngapml
Global active subscriber limit	NATLIM	GSLIM	3 error	ngsl
Subscriber-specific address and port mapping limit	NATLIM	SAPMLIM	5 notice	nsapml
Pending fragment limit	NATLIM	FRAG	4 warning	nfpkt

Table 1: Recommended MSGID Encodings and Default Severity 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. The PARAM-NAMEs, brief

descriptions, and encoding are listed in Table 2, with reference to the general and special case encoding rules which follow.

PARAM-NAME	Description	Encoding
	Miscellaneous	
NATINST	NAT instance identifier	Text
TRIG	Trigger for event	Special case
	Subscriber-identifying information	
SSUBIX	Source subscriber index	32-bit field
SIFIX	Source subscriber ingress interface index list	Special case
SVLAN	Source subscriber ingress VLAN index	32-bit field
SVPN	Source subscriber ingress VPN Id	Special case
SV6ENC	Source subscriber ingress RFC6333 encapsulating IPv6 address	IPv6 address
DSUBIX	Destination subscriber index	32-bit field
DIFIX	Destination subscriber ingress interface index list	Special case
DVLAN	Destination subscriber ingress VLAN index	32-bit field
DVPN	Destination subscriber ingress VPN Id	Special case
DV6ENC	Destination subscriber ingress RFC6333 encapsulating IPv6 address	IPv6 address
	Internal packet description	
IRLM	Internal realm	Text
IATYP	Internal IP address type	"IPv4" or "IPv6"
ISADDR	Internal source IP address value	IPv4 or IPv6 address
ISPORT	Internal source port or ICMP identifier value	16-bit field
IDADDR	Internal destination IP address value	IPv4 or IPv6 address

IDPORT	Internal destination port or ICMP identifier value	16-bit field
PROTO	Protocol identifier (from the IANA Assigned Internet Protocol Numbers registry)	8-bit field
	External (mapped) packet description	
XRLM	External realm	Text
XATYP	External IP address type	"IPv4" or "IPv6"
XSADDR	External source IP address value	IPv4 or IPv6 address
XSPORT	External source port or ICMP identifier value	16-bit field
XDADDR	External destination IP address value	IPv4 or IPv6 address
XDPORT	External destination port or ICMP identifier value	16-bit field
	Port range description	
PORTMN	Port range lowest value	16-bit field
PORTMX	Port range highest value	16-bit field
	Values related to thresholds	
POOLID	Address pool identifier	32-bit field
POOLHW	Address pool high water mark threshold	Unsigned decimal
POOLID	Address pool low water mark threshold	Unsigned decimal
GAMCNT	Current global number of address mappings	Unsigned decimal
GAPMCNT	Current global number of address and port mappings	Unsigned decimal
SAPMCNT	Current subscriber-specific number of address and port mappings	Unsigned decimal
	Specific incoming packet description	
PSRLM	Packet source realm	Text
PATYP	Packet IP address type	"IPv4" or

		"IPv6"
PSADDR	Packet source IP address	IPv4 or IPv6 address
PDADDR	Packet destination IP address	IPv4 or IPv6 address

Table 2: Parameters Used In NAT-Related Log Reports

5.2.1. General Encoding Rules

All fields MUST be encoded as 7-bit US ASCII [US-ASCII].

Complete IPv6 addresses MUST be presented according to the rules specified in Sections 4 and 5 of [RFC5952], without a succeeding prefix length. The Section 5 rules MUST NOT be applied unless the address can be distinguished as having an IPv4 address embedded in the lower 32 bits solely from the IPv6 prefix portion (e.g., based on well-known prefix, flag), without external information. In such cases, the IPv6 prefix portion MUST be presented according to the Section 4 rules. Stand-alone IPv6 prefixes (i.e., outside of special addresses) MUST be presented according to the Section 4 rules, with the slash character (/) appended, followed by a decimal value with leading zeroes suppressed, giving the prefix length (0 to 127) in bits.

Similarly, complete IPv4 addresses MUST be presented in dotted decimal format, with no succeeding prefix length. IPv4 prefixes MUST be presented as if they were full addresses, with the slash character (/) appended, followed by a decimal value with leading zeroes suppressed, giving the prefix length (0 to 31) in bits.

N-bit fields and unsigned decimals are both presented as unsigned decimal integers with no leading zeroes.

5.2.2. Special Cases

Three special cases are identified in Table 2: encoding of the interface index list (PARAM-NAMEs SIFIX and DIFIX), encoding of the VPN identifier (PARAM-NAMEs SVPN and DVPN), and encoding of the trigger for resource allocation events (PARAM-NAME TRIG).

The interface index list is presented as a series of individual interface indexes separated by commas, e.g., SIFIX="5,15". Each individual interface index is presented as a 32-bit field (i.e., as an unsigned decimal integer with no leading zeroes).

The VPN Identifier is standardized in [RFC2685], and consists of a three octet VPN Authority (Organizationally Unique Identifier, OUI) followed by a four octet VPN index identifying the VPN according to OUI. For SYSLOG, the OUI portion is presented as a string of six hexadecimal digits in lower case. The VPN index is presented as a 32-bit field. A colon (:) is used to separate the OUI from the succeeding index value. The OUI and separator MAY be omitted. If so, the applicable OUI is the default value for the NAT instance.

The trigger is an enumeration of text values which were not spelled out in the table itself for lack of space. The possible values for TRIG are:

"OPKT": outgoing packet received at NAT.

"IPKT": incoming packet received at NAT.

"ADMIN": administrative action.

"APMDEL": deletion of the underlying address and port mapping.

"AMDEL": deletion of the underlying address mapping.

"AUTO": autonomous action of the NAT.

The values applicable for any specific event are a subset of this list and are spelled out for each event in Section 5.3.

5.2.3. Relationship To Objects In the NAT MIB

Table 3 lists the parameters in the same order as Table 2 and relates each parameter to its corresponding object in the NAT MIB.

PARAM-NAME	Related MIB Object(s)
Miscellaneous	
NATINST	natInstanceAlias in natInstanceTable
TRIG	None
Subscriber-identifying	
information	
SSUBIX	natSubscriberIndex in
	natSubscribersTable
SIFIX	natSubsInterfaceIndex in
	natSubsInterfaceIdentifierTable

SVLAN	natSubscriberVlanIdentifier in natSubscribersTable
SVPN	natSubscriberVpnIdentifier in natSubscribersTable
SV6ENC	natSubscriberIPEncapsIdType and natSubscriberIPEncapsIdAddr in natSubscribersTable
DSUBIX	natSubscriberIndex in natSubscribersTable
DIFIX	natSubsInterfaceIndex in natSubsInterfaceIdentifierTable
DVLAN	natSubscriberVlanIdentifier in natSubscribersTable
DVPN	natSubscriberVpnIdentifier in natSubscribersTable
DV6ENC	natSubscriberIPEncapsIdType and natSubscriberIPEncapsIdAddr in natSubscribersTable
Internal packet description	
IRLM	natSubscriberRealm in natSubscribersTable
IATYP	natMapIntAddrIntType in natMapIntAddrTable or natMappingIntAddressType in natMappingTable
ISADDR	natMapIntAddrInt in natMapIntAddrTable or natMappingIntAddress in natMappingTable
ISPORT	natMappingIntPort in natMappingTable
IDADDR	None
IDPORT	None
PROTO	natMappingProto in natMappingTable
External (mapped) packet description	
XRLM	natPoolRealm in natPoolTable
XATYP	natMapIntAddrExtType in natMapIntAddrTable or natMappingExtAddressType in natMappingTable
XSADDR	natMapIntAddrExt in natMapIntAddrTable or natMappingExtAddress in natMappingTable
XSPORT	natMappingExtPort in natMappingTable

XDADDR	None
XDPORT	None
Port range description	
PORTMN	None
PORTMX	None
Values related to thresholds	
POOLID	natPoolIndex in natPoolTable
POOLHW	natPoolWatermarkHigh in natPoolTable
POOLLW	natPoolWatermarkLow in natPoolTable
GAMCNT	natAddressMappingCreations - natAddressMappingRemovals in natCountersTable
GAPMCNT	natMappingCreations - natMappingRemovals in natCountersTable
SAPMCNT	natSubscriberMappingCreations - natSubscriberMappingRemovals in natSubscribersTable
Specific incoming packet description	
PSRLM	natSubscriberRealm in natSubscribersTable in the case of a packet originated by an identifiable subscriber
PATYP	None
PSADDR	None
PDADDR	None

Table 3: Relationship of Parameters To Objects In the NAT MIB

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. Encoding of Events Relating To Allocation Of Resources To Hosts

As indicated in Section 5.1, the event reports specified in this section MUST have APP-NAME="NAT" in the message header.

5.3.1.1. NAT Address Mapping Creation and Deletion

As shown in Table 1:

- NAT address mapping creation event is indicated by MSGID set to "AMADD";

- NAT address mapping deletion event is indicated by MSGID set to "AMDEL".

For both events, the associated SD-ELEMENT is tagged by SD-ID "namap". The contents of the namap SD-ELEMENT are shown in Table 4. The requirements for these contents are derived from the description in Section 3.1.1.

Description	PARAM-NAME	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Source subscriber index	SSUBIX	MANDATORY
Additional source subscriber classifier value as recognized at the ingress to the internal realm	One of SIFIX, SVLAN, SVPN, or SV6ENC	CONDITIONAL
Internal realm	IRLM	CONDITIONAL
Internal address type	IATYP	MANDATORY
Internal source IP address	ISADDR	MANDATORY
External realm	XRLM	CONDITIONAL
External address type	XATYP	MANDATORY
External source IP address	XSADDR	MANDATORY
Trigger for address mapping creation or deletion	TRIG	OPTIONAL

Table 4: Contents Of the SD-ELEMENT Section For Logging the Address Mapping Creation and Deletion Events

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

- One of SIFIX, SVLAN, SVPN, or SV6ENC REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.

- IRLM or XRLM REQUIRED if not the default internal or external realm, respectively, else MAY appear.

For the AMADD event type (MSGID), TRIG can take on the values "OPKT" or "ADMIN". For the AMDEL event type, TRIG can take on the values "ADMIN" or "AUTO".

Example: DS-Lite AFTR. One NAT instance. Multiple internal IPv4 realms containing the subscribers, divided by higher-level IPv6 prefix (details unnecessary). One default global IPv4 external realm. Intra-subscriber sessions use mappings into this realm.

Subscriber A in realm Internal05 sends an outgoing packet, causing the creation of an address mapping from the DS-Lite well-known address 192.0.0.2 to the global IPv4 address 198.51.100.127. Subscriber A's encapsulating IPv6 tunnel address is 2001:db8:a5e6:39b0:bd6a:35ad:1d33:6df6.

The event report for the address mapping creation is as follows (line folded into several for presentation):

```
<142>1 2013-05-07T22:14:15.03487Z record.example.net NAT 5063
AMADD [namap SSUBIX="489321"
SV6ENC="2001:db8:a5e6:3900:bd6a:35ad:1d33:6df6" IRLM="Internal05"
IATYP="IPv4" ISADDR="192.0.0.2" XATYP="IPv4"
XSADDR="198.51.100.127"
TRIG="OPKT"]
```

Character count is about 240.

5.3.1.2. NAT Address and Port Mapping Creation and Deletion

As shown in Table 1:

- NAT address and port mapping creation event is indicated by MSGID set to "APMADD";

- NAT mapping deletion event is indicated by MSGID set to "APMDEL".

For both events, the associated SD-ELEMENT is tagged by SD-ID "napmap". The contents of the nmap SD-ELEMENT are shown in Table 5. The requirements for these contents are derived from the description in Section 3.1.2.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Source subscriber index	SSUBIX	MANDATORY
Additional source subscriber classifier value as recognized at the ingress to the internal realm	One of SIFIX, SVLAN, SVPN, or SV6ENC	CONDITIONAL
Internal realm	IRLM	CONDITIONAL
Internal address type	IATYP	MANDATORY
Internal source IP address	ISADDR	MANDATORY
Internal source port or ICMP identifier	ISPORT	MANDATORY
External realm	XRLM	CONDITIONAL
External address type	XATYP	MANDATORY
External source IP address	XSADDR	MANDATORY
External source port or ICMP identifier	XSPORT	MANDATORY
Protocol identifier	PROTO	MANDATORY
Trigger for address and port mapping creation or deletion	TRIG	OPTIONAL

Table 5: Contents Of the SD-ELEMENT Section For Logging the mapping Creation and Deletion Events

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

- One of SIFIX, SVLAN, SVPN, or SV6ENC REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.

- IRLM or XRLM REQUIRED if not the default internal or external realm, respectively, else MAY appear.

For the APMADD event type (MSGID), TRIG can take on the values "OPKT", IPKT", or "ADMIN".

Note: it is not clear how the internal source port is selected if an address and port mapping is triggered by an incoming TCP packet. The NAT could select one based on its knowledge of subscriber port usage, but this knowledge may be incomplete. Some type of negotiation may be necessary, or else TCP address and port mappings can only be triggered by outbound packets as in the example below.

For the APMDEL event type, TRIG can take on the values "ADMIN", "AMDEL", or "AUTO".

Example: The triggering outgoing packet in the previous case was a TCP packet with internal source port 49178. As well as triggering the creation of an address mapping, the packet triggers the creation of an address and port mapping between that port and an external source port 6803. The corresponding mapping creation report would look like this:

```
<142>1 2013-05-07T22:14:15.03487Z record.example.net NAT 5063
APMADD [napmap SSUBIX="489321"
SV6ENC="2001:db8:a5e6:3900:bd6a:35ad:1d33:6df6" IRLM="Internal05"
IATYP="IPv4" ISADDR="192.0.0.2" ISPORT="49178"
XATYP="IPv4" XSADDR="198.51.100.127" XSPORT=6803"
PROTO="6" TRIG="OPKT"]
```

Character count is about 280.

5.3.1.3. NAT Session Creation and Deletion

As shown in Table 1:

- NAT session creation event is indicated by MSGID set to "SADD";

- NAT session deletion event is indicated by MSGID set to "SDEL".

For both events, the associated SD-ELEMENT is tagged by SD-ID "nsess". The contents of the nsess SD-ELEMENT are shown in Table 6. The requirements for these contents are derived from the description in Section 3.1.3.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Source subscriber index	SSUBIX	MANDATORY
Additional source subscriber classifier value as recognized at the ingress to the internal realm	One of SIFIX, SVLAN, SVPN, or SV6ENC	CONDITIONAL
Internal realm	IRLM	CONDITIONAL
Internal address type	IATYP	MANDATORY
Internal source IP address	ISADDR	MANDATORY
Internal source port or ICMP identifier	ISPORT	MANDATORY
External realm	XRLM	CONDITIONAL
External address type	XATYP	MANDATORY
External source IP address	XSADDR	MANDATORY
External source port or ICMP identifier	XSPORT	MANDATORY
Protocol identifier	PROTO	MANDATORY
Internal destination IP address	IDADDR	CONDITIONAL
Internal destination port or ICMP identifier	IDPORT	CONDITIONAL
Destination subscriber index	DSUBIX	CONDITIONAL
Additional destination subscriber classifier value as recognized at the ingress to the external realm	One of DIFIX, DVLAN, DVPN, or DV6ENC	CONDITIONAL
External destination IP address	XDADDR	CONDITIONAL
External destination port or ICMP identifier	XDPORT	CONDITIONAL
Trigger for session creation or deletion	TRIG	OPTIONAL

Table 6: Contents Of the SD-ELEMENT Section For Logging the Session Creation and Deletion Events

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

- One of SIFIX, SVLAN, SVPN, or SV6ENC REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.

- IRLM or XRLM REQUIRED if not the default internal or external realm, respectively, else MAY appear.

- IDADDR and IDPORT REQUIRED if destination logging is enabled and these need to be remapped to external destination address and port. Otherwise, if destination logging is disabled, they MUST NOT appear, and if destination logging is enabled, they SHOULD NOT appear because of redundancy.

- DSUBIX REQUIRED if destination logging is enabled and the destination is a subscriber served by the NAT, else MUST NOT appear.

- One of DIFIX, DVLAN, DVPN, or DV6ENC REQUIRED if destination logging is enabled and the destination is a subscriber served by the NAT and the external destination address is not enough to identify the external destination subscriber unambiguously, else MUST NOT appear.

- XDADDR and XDPORT REQUIRED if destination logging is enabled, else MUST NOT appear.

For the SADD event type (MSGID), TRIG can take on the values "OPKT", IPKT", or "ADMIN". For the SDEL event type, TRIG can take on the values "ADMIN", "MDEL", or "AUTO".

Example: destination logging is enabled. The outgoing packet that triggered the address and port mapping in the previous section was sent to 192.0.2.57 port 80. The session creation event report appears as follows:

```
<142>1 2013-05-07T22:14:15.03487Z record.example.net NAT 5063
SESSADD [nsess SSUBIX="489321"
SV6ENC="2001:db8:a5e6:3900:bd6a:35ad:1d33:6df6" IRLM="Internal05"
IATYP="IPv4" ISADDR="192.0.0.2" ISPORT="49178"
XATYP="IPv4" XSADDR="198.51.100.127" XSPORT=6803"
PROTO="6" XDADDR="192.0.2.57" XDPORT="80" TRIG="OPKT"]
```

Character count is about 310.

5.3.1.4. Port Range Allocation and Deallocation

As shown in Table 1:

- Port range allocation event is indicated by MSGID set to "PTADD";

- Port range deallocation event is indicated by MSGID set to "PTDEL".

For both events, the associated SD-ELEMENT is tagged by SD-ID "nprng". The contents of the npset SD-ELEMENT are shown in Table 7.

The requirements for these contents are derived from the description in Section 3.1.4.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Source subscriber index	SSUBIX	MANDATORY
Additional source subscriber classifier value as recognized at the ingress to the internal realm	One of SIFIX, SVLAN, SVPN, or SV6ENC	CONDITIONAL
Internal realm	IRLM	CONDITIONAL
Internal address type	IATYP	MANDATORY
Internal source IP address	ISADDR	MANDATORY
External realm	XRLM	CONDITIONAL
External address type	XATYP	MANDATORY
External source IP address	XSADDR	MANDATORY
Port range lowest value	PORTMN	MANDATORY
Port range highest value	PORTMX	MANDATORY
Trigger for port range allocation or deallocation	TRIG	OPTIONAL

Table 7: Contents Of the SD-ELEMENT Section For Logging the Port Set Allocation and Deallocation Events

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

- One of SIFIX, SVLAN, SVPN, or SV6ENC REQUIRED if the internal source IP address is not enough to identify the subscriber unambiguously, else MUST NOT appear.

- IRLM or XRLM REQUIRED if not the default internal or external realm, respectively, else MAY appear.

For the PTADD event type (MSGID), TRIG can take on the values "OPKT", "IPKT", "ADMIN", or "AUTO". For the PTDEL event type, TRIG can take on the values "ADMIN" or "AUTO".

Consider an example where the range 1024-1535 is allocated to the address mapping on which the example in Section 5.3.1.1 is based. The corresponding port range allocation report would look like this:

```
<142>1 2013-05-07T22:14:15.03487Z record.example.net NAT 5063
PTADD [nprng SSUBIX="489321"
```

```
SV6ENC="2001:db8:a5e6:3900:bd6a:35ad:1d33:6df6" IRLM="Internal05"
IATYP="IPv4" ISADDR="192.0.0.2" XATYP="IPv4"
XSADDR="198.51.100.127"
PORTMN="1024" PORTMX="1535" TRIG="OPKT"]
```

Character count is about 270.

5.3.2. Encoding of Threshold Events

As indicated in Section 5.1, the event reports specified in this section MUST have APP-NAME="NATTHR" in the SYSLOG message header.

5.3.2.1. NAT Address Pool High- and Low-Water-Mark Threshold Events

As shown in Table 1:

- NAT address pool high-water-mark threshold event is indicated by MSGID set to "POOLHT";
- NAT address pool low-water-mark threshold event is indicated by MSGID set to "POOLLT".

For both events, the associated SD-ELEMENT is tagged by SD-ID "npool". The contents of the npool SD-ELEMENT are shown in Table 8. The requirements for these contents are derived from the description in Section 3.2.1.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Pool identifier	POOLID	MANDATORY
The threshold value set by the administrator	POOLHW or POOLLW as applicable	CONDITIONAL

Table 8: Contents Of the SD-ELEMENT Section For Logging the Address Pool High- and Low-Water-Mark Threshold Events

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.
- POOLHW REQUIRED for high-water-mark event, else MUST NOT appear.
- POOLLW REQUIRED for low-water-mark event, else MUST NOT appear.

Example, assuming a high-water-mark threshold of 80% aggregate address-port utilization::

```
<132>1 2013-08-15T09:15:16.08716Z record.example.net NATTHR 5025
POOLHT [npool POOLID="13" POOLHW="80"]
```

Character count is about 105.

5.3.2.2. Global Address Mapping High-Water-Mark Threshold Exceeded

As shown in Table 1:

- Global address mapping high-water-mark threshold event is indicated by MSGID set to "GAMHT"; and
- the associated SD-ELEMENT is tagged by SD-ID "ngamht".

The contents of the ngamht SD-ELEMENT are shown in Table 9. The requirements for these contents are derived from the description in Section 3.2.2.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Current number of active address mappings	GAMCNT	MANDATORY

Table 9: Contents Of the SD-ELEMENT Section For Logging the Global Address Map High-Water-Mark Threshold Event

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

Example, assuming a threshold was set to 690000, already exceeded. As a result, prior events of this type were detected and logged, unless they were suppressed by the sort of controls discussed in Section 6.

```
<132>1 2013-08-15T09:15:16.08716Z record.example.net NATTHR 5025
GAMHT [ngamht GAMCNT="690015"]
```

Character count is about 95.

5.3.2.3. Global Address and Port Mapping High-Water-Mark Threshold Event

As shown in Table 1:

- Global address and port mapping high-water-mark threshold event is indicated by MSGID set to "GAPMHT"; and

- the associated SD-ELEMENT is tagged by SD-ID "ngapmht".

The contents of the ngmht SD-ELEMENT are shown in Table 10. The requirements for these contents are derived from the description in Section 3.2.3.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Current global number of address and port mappings	GAPMCNT	MANDATORY

Table 10: Contents Of the SD-ELEMENT Section For Logging the Global Address and Port Mapping High-Water-Mark Threshold Event

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

Example: suppose the threshold was set to 2000000, so it has already been exceeded. As in the previous section, prior events of this type were detected and logged, unless they were suppressed by the sort of controls discussed in Section 6.

```
<132>1 2013-08-15T09:15:16.08716Z record.example.net NATTHR 5025
GAPMHT [ngapmht GAPMCNT="2000023"]
```

Character count is about 100.

5.3.2.4. Subscriber-Specific Address and Port Mapping High-Water-Mark Threshold Event

As shown in Table 1:

- Subscriber-specific address and port mapping high-water-mark threshold event is indicated by MSGID set to "SAPMHT"; and

o the associated SD-ELEMENT is tagged by SD-ID "nsapmht".

The contents of the nsapmht SD-ELEMENT are shown in Table 11. The requirements for these contents are derived from the description in Section 3.2.4.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Index of the affected subscriber	SSUBIX	MANDATORY
Current number of address and port mappings for this subscriber	SAPMCNT	MANDATORY

Table 11: Contents Of the SD-ELEMENT Section For Logging the Subscriber-Specific Address and Port Mapping High-Water-Mark Threshold Event

Conditions:

o NATINST REQUIRED if device supports more than one instance, else MAY appear.

Example: suppose the threshold was set to 1500 and the number of mappings for this subscriber has been increasing. Then this is the first threshold-exceeded event detected of what could possibly be a series of such events until subscriber consumption of outgoing ports drops below threshold again.

```
<133>1 2013-08-15T09:15:16.08853Z record.example.net NATTHR 5025
SAPMHT [nsapmht SSUBIX="489321" SAPMCNT="1501"]
```

Character count is about 115.

5.3.3. Encoding of Limit Events

As indicated in Section 5.1, the event reports specified in this section MUST have APP-NAME="NATLIM" in the SYSLOG message header.

5.3.3.1. Global Address Mapping Limit Exceeded

As shown in Table 1:

o Global address mapping limit exceeded event is indicated by MSGID set to "GAMLIM"; and

o the associated SD-ELEMENT is tagged by SD-ID "ngaml".

The contents of the ngaml SD-ELEMENT are shown in Table 12. The requirements for these contents are derived from the description in Section 3.3.1.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Index of the affected subscriber	SSUBIX	MANDATORY

Table 12: Contents Of the SD-ELEMENT Section For Logging the Global Address Map Limit Exceeded Event

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

Example:

```
<131>1 2013-08-15T09:15:16.08716Z record.example.net NATLIM 5025
GAMLIM [ngaml NATINST="VRF-Cust-X" SSUBIX="278067"]
```

Character count is about 115.

5.3.3.2. Global Address and Port Mapping Limit Exceeded

As shown in Table 1:

- Global adress and port mapping limit exceeded event is indicated by MSGID set to "GAPMLIM"; and

- the associated SD-ELEMENT is tagged by SD-ID "ngapml".

The contents of the ngapml SD-ELEMENT are shown in Table 13. The requirements for these contents are derived from the description in Section 3.3.2.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Index of the internal subscriber	SSUBIX	CONDITIONAL
Index of the external subscriber	DSUBIX	CONDITIONAL
Source realm of the triggering packet	PSRLM	MANDATORY
Incoming packet header IP address type	PATYP	CONDITIONAL
Incoming packet source IP address	PSADDR	CONDITIONAL

Table 13: Contents Of the SD-ELEMENT Section For Logging the Global Address and Port Mapping Limit Exceeded Event

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

- SSUBIX REQUIRED if the mapping was triggered by a packet outgoing from the internal to the external realm, else MUST NOT appear.

- DSUBIX is REQUIRED if the mapping was triggered by a packet incoming from a subscriber served by the NAT and located in the external realm (i.e., using an address mapping created previously by the internal subscriber), else MUST NOT appear.

- PATYP and PSADDR from the initiating packet are REQUIRED if the mapping was triggered by a packet incoming from a purely external realm (i.e., using an address mapping created previously by the internal subscriber), else MAY appear.

Example: limit event triggered by a packet coming from 192.0.2.57 in realm "externv4".

```
<131>1 2013-08-15T09:15:16.08716Z record.example.net NATLIM 5025
GAPMLIM [ngapml NATINST="VRF-Cust-X" PSRLM="externv4"
PATYP="IPv4" PSADDR="192.0.2.57"]
```

Character count is about 150.

5.3.3.3. Global Limit On Number of Active Hosts Exceeded

As shown in Table 1:

- Global active hosts limit exceeded event is indicated by MSGID set to "GSLIM"; and

o the associated SD-ELEMENT is tagged by SD-ID "ngsl".

The contents of the ngsl SD-ELEMENT are shown in Table 14. The requirements for these contents are derived from the description in Section 3.3.3.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Index of the affected subscriber	SSUBIX	MANDATORY

Table 14: Contents Of the SD-ELEMENT Section For Logging the Global Active Host Limit Exceeded Event

Conditions:

o NATINST REQUIRED if device supports more than one instance, else MAY appear.

An example would look exactly like that in Section 5.3.3.1 with the substitution of GSLIM for GAMLIM and ngsl for ngaml.

5.3.3.4. Subscriber-Specific Limit On Number of Address and Port Mappings Exceeded

As shown in Table 1:

o Subscriber-specific mapping limit exceeded event is indicated by MSGID set to "SMLIM"; and

o the associated SD-ELEMENT is tagged by SD-ID "nsml".

The contents of the nsml SD-ELEMENT are shown in Table 15. The requirements for these contents are derived from the description in Section 3.3.4.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Index of the affected subscriber	SSUBIX	MANDATORY

Table 15: Contents Of the SD-ELEMENT Section For Logging the Subscriber-Specific Mapping Limit Exceeded Event

Conditions:

- NATINST REQUIRED if device supports more than one instance, else MAY appear.

An example would look exactly like that in Section 5.3.3.1 with the substitution of SAPMLIM for GAMLIM and nsapml for ngaml.

5.3.3.5. Pending Fragment Limit Exceeded

As shown in Table 1:

- Pending fragment limit exceeded event is indicated by MSGID set to "FRAG"; and

- the associated SD-ELEMENT is tagged by SD-ID "nfpkt".

The contents of the nfpkt SD-ELEMENT are shown in Table 16. The requirements for these contents are derived from the description in Section 3.3.5.

PARAM-NAME	Description	Requirement
NAT instance identifier	NATINST	CONDITIONAL
Source realm of the packet	PSRLM	MANDATORY
Packet header IP address type	PATYP	MANDATORY
Packet source IP address	PSADDR	MANDATORY
Packet destination IP address	PDADDR	MANDATORY
Source subscriber index	SSUBIX	CONDITIONAL

Table 16: Contents Of the SD-ELEMENT Section For Logging the Pending Fragment Limit Exceeded Event

Conditions:

- NAT instance identifier REQUIRED if device supports more than one instance, else MAY appear.

- Source subscriber index REQUIRED if the source of the packet is a subscriber served by the NAT and can be determined, else MUST NOT appear.

Example: assuming the packet passing the limit came from an internal host and was dropped as a result of the limit.

```
<132>1 2013-08-15T09:15:16.08Z record.example.net NATLIM 5025
```

```
FRAG [nfpkt PSRLM="DsLite-089" PATYP="IPv4" PSADDR="192.0.0.2"
PDADDR="203.0.113.26" SSUBIX="32791"]
```

Character count is about 160.

6. 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 fulfil these requirements is out of scope of the present document. Subsequent sub-sections discuss management issues related to specific event report types. The identifiers PRI, APP-NAME, and MSGID used below refer to fields in the SYSLOG header [RFC5424]

6.1. General Requirements For Control Of Logging

This document assumes that any implementation provides the following capabilities, discussed in more detail below:

- ability to configure the PRI value of each event report type at the granularity of (APP-NAME, MSGID) combination;

- ability at each collector to determine that event reports that it should have received have been lost. The required granularity is at least at the level of PRI and may be finer for some event types.

- ability to configure criteria to automatically suppress the generation of event reports while the criteria are met, at the granularity of (APP-NAME, MSGID) combination.

6.1.1. Configuration of PRI Value

The PRI value is composed of two numbers, the Facility value and the Severity. It may be used at the origin for selecting logs to streams being dispatched to different collectors, and in applications beyond the collectors to prioritize display of logs to operators. The event reports in this document have been structured such that the Severity level varies between event types as represented by (APP-NAME, MSGID) combination. As an extreme example, the address pool high-water-mark threshold event (APP-NAME="NATTHR", MSGID="POOLHT") is obviously more urgent than the low-water-mark threshold event (APP-NAME="NATTHR", MSGID="POOLLT").

To some extent, this document tries to simplify message routing by making a general distinction between event types recording the

allocation of resources to hosts (with APP-NAME="NAT") and events of interest to operations and maintenance (with APP-NAME="NATTHR" and APP-NAME="NATLIM"). The need to provide different Severity levels for different event types remains.

6.1.2. Ability For Each Collector To Detect Lost Event Reports

Operators have a need to know when a given collector has not received all of the event reports it should have. It probably does not matter if less-important events are tracked at the granularity of event type (APP-NAME, MSGID combination), by APP-NAME, or just by PRI value.

The event types defined in this document relating to allocation of resources to hosts are a special case. Regulatory requirements or the possibility that such reports might be introduced into court in cases such as abuse impose a requirement that the record of allocations to a particular host be complete. This requirement is important enough to be stated in the Security Considerations section (Section 7), where the implementation of signed SYSLOG messages [RFC5848], which also provides message sequencing, is mandated as part of this specification.

In deploying [RFC5848], the operator needs to decide the level of granularity of tracking, whether it should be over the whole set of reports covered by APP-NAME="NAT" or at a finer level. This judgement has to be tempered by local circumstances. One point to note is that since both creations/allocations and deletions/deallocations are recorded, a certain amount of redundancy is available in the reports being generated. However, without both the creation and deletion timestamps, there is no definitive evidence of the specific period of time during which the resources concerned were allocated to a specific host.

6.1.3. Ability To Rate Limit Or Disable Event Reports

The event report types specified with APP-NAME="NATTHR" and APP-NAME="NATLIM" all relate to thresholds or limits. By their nature, events of this sort will come in bursts. The threshold or limit will be hit, the resource concerned will remain busy for a period, then pressure on the resource will ease. Depending on the resource, possibly hundreds of instances of the event concerned will be detected during a single busy period.

Where repeated events involve the same resource, it makes little sense to report all of them, since the NAT MIB counters provide the necessary information more succinctly. On the other hand, it can be useful to know that the fragmentation limit, for instance, is being hit by successive packets from the same source address.

As a result of these considerations, this document requires that implementations MUST provide means to configure limits on the rate at which event reports of a given type (APP-NAME, MSGID combination) are generated. It is RECOMMENDED that it be possible to specify two values per (APP-NAME, MSGID) combination:

- minimum time between initial instances of a given event report type;

- maximum number of instances of the event report to generate per busy period.

Regardless of the detailed method the implementation provides for specifying the rate limiting of individual event report types, all implementations MUST allow the operator to indicate through configuration that a given event report type is to be completely disabled. This is particularly required to disable logging of either session or mapping creations and deletions when not required (see discussion in Section 3.1.2). It is also required when the operator prefers to receive threshold event notifications via SNMP rather than SYSLOG.

The ability to rate limit or disable event reports MUST NOT interfere with the requirement to detect lost messages. This has implications for any sequence numbering used for that purpose. It is RECOMMENDED in any event that the implementation provide MIB counters of numbers of messages not generated due to rate limiting by event type supported. If this is done, counters for disabled event report types SHOULD NOT be incremented, since that could require keeping unnecessary additional state.

6.2. Setting Limits and Thresholds

The "NATTHR" and "NATLIM" events specified in this document depend on the thresholds and limits configured in the NAT MIB [I-D.behave-NAT-MIB]. The limits have to do with policy in some cases (e.g., most especially the subscriber-specific limits), but generally depend on the implementation and the device in which it is deployed.

The purpose of high-water-mark thresholds is, of course, to give sufficient advance warning that utilization of a particular resource is approaching its limit, so that appropriate provisioning or reconfiguration action can be undertaken to preserve target service levels on the NAT device. Thus the following general principles apply:

- A high-water-mark threshold should be derived as a percentage of the relevant limit.

- The more quickly that utilization of a given resource can build up, the lower the threshold must be to provide an adequate response time.

- Some limits are more important than others in terms of their effect on overall service levels provided by the NAT device. To focus attention on the more important limits, their corresponding thresholds should be set lower than those for less-important limits, all other things being equal.

In practice, thresholds will require tuning to fit the particular characteristics of the NAT device and its users.

The setting of the high-water-mark-thresholds for address pools (Section 3.2.1) poses additional challenges. The problem is that the bottleneck for port availability will generally be a single protocol, which may vary from one time to another. However, the threshold is based on overall port utilization. If port usage is such that one protocol generally predominates, the required threshold value has to be lower than if usage is more balanced between protocols. Clearly the appropriate threshold value depends on the characteristics of the traffic handled by the particular address pool concerned.

Pooling behaviour adds another factor for consideration. With a pooling behaviour of "arbitrary" [RFC4787], port utilization for the bottleneck protocol can be quite high before service levels offered by the pool are in danger. On the other hand, with a pooling behaviour of "paired", possible utilization levels will be much lower because typically a number of port values will be reserved to each address mapping and only some of those will be in use on the average. The difference between "arbitrary" and "paired" utilization for a given level of service may be quite dramatic.

6.3. Other Management Requirements

The identification of internal realms is contingent on the the existence and applicability of default internal and external realms. If the implementation is capable of supporting more than one internal or external realm, it MUST provide the means for the operator to specify which realm is the default internal and/or external realm, as the case may be.

7. Security Considerations

When logs are being recorded for regulatory reasons or as potential evidence in abuse cases, preservation of their integrity and authentication of their origin is essential. To achieve this result, signed SYSLOG messages [RFC5848] MUST be implemented as part of this specification. It is RECOMMENDED that the operator deploy [RFC5848] where local requirements on integrity and authentication of origin are stringent. In conjunction with [RFC5848] and as recommended in Section 3 of that document, TLS transport as specified in [RFC5425] SHOULD be used between the origin and the collector(s) and MUST be implemented. Section 5.2.1 of [RFC5848] specifies the minimum support for Key Blob Type that must be provided by implementations of that specification.

Access to the logs defined in Section 3.1 and Section 5.3.1 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 user or user location. It is therefore RECOMMENDED that these logs be transported securely, using [RFC5425], for example, even if [RFC5848] is not deployed, that they be stored securely at the collector, and that access to them at the collector and in applications be tightly controlled.

The logs defined in Section 3.2 and Section 3.3 are less sensitive in general, but since many of them contain the subscriber identifier, they could be used to get some sense of subscriber activity. The fragmentation limit event provides actual packet header contents. Operators SHOULD at the least deploy secure transport to ensure that this information is not misused.

8. 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.

Some PARAM-NAMES appear under more than one SD-ID in Table 17. Formally, 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 if the parameters have the same meaning in both event reports.

While a number of parameters are marked CONDITIONAL in the body of this document, the SYSLOG registry provides only for MANDATORY and OPTIONAL parameters. All CONDITIONAL parameters have been placed in the OPTIONAL category in Table 17.

Structured Data ID	Structured Data Parameter	Required or Optional	Reference
namap		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
	SIFIX	OPTIONAL	RFCxxxx
	SVLAN	OPTIONAL	RFCxxxx
	SVPN	OPTIONAL	RFCxxxx
	SV6ENC	OPTIONAL	RFCxxxx
	IRLM	OPTIONAL	RFCxxxx
	IATYP	MANDATORY	RFCxxxx
	ISADDR	MANDATORY	RFCxxxx
	XRLM	OPTIONAL	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XSADDR	MANDATORY	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
napmap		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
	SIFIX	OPTIONAL	RFCxxxx
	SVLAN	OPTIONAL	RFCxxxx
	SVPN	OPTIONAL	RFCxxxx
	SV6ENC	OPTIONAL	RFCxxxx
	IRLM	OPTIONAL	RFCxxxx
	IATYP	MANDATORY	RFCxxxx
	ISADDR	MANDATORY	RFCxxxx
	ISPORT	MANDATORY	RFCxxxx
	XRLM	OPTIONAL	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XSADDR	MANDATORY	RFCxxxx
	XSPORT	MANDATORY	RFCxxxx
	PROTO	MANDATORY	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
nsess		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
	SIFIX	OPTIONAL	RFCxxxx
	SVLAN	OPTIONAL	RFCxxxx
	SVPN	OPTIONAL	RFCxxxx

	SV6ENC	OPTIONAL	RFCxxxx
	IRLM	OPTIONAL	RFCxxxx
	IATYP	MANDATORY	RFCxxxx
	ISADDR	MANDATORY	RFCxxxx
	ISPORT	MANDATORY	RFCxxxx
	XRLM	OPTIONAL	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XSADDR	MANDATORY	RFCxxxx
	XSPORT	MANDATORY	RFCxxxx
	PROTO	MANDATORY	RFCxxxx
	IDADDR	OPTIONAL	RFCxxxx
	IDPORT	OPTIONAL	RFCxxxx
	DSUBIX	OPTIONAL	RFCxxxx
	DIFIX	OPTIONAL	RFCxxxx
	DVLAN	OPTIONAL	RFCxxxx
	DVPN	OPTIONAL	RFCxxxx
	DV6ENC	OPTIONAL	RFCxxxx
	XDADDR	OPTIONAL	RFCxxxx
	XDPORT	OPTIONAL	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
nprng		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
	SIFIX	OPTIONAL	RFCxxxx
	SVLAN	OPTIONAL	RFCxxxx
	SVPN	OPTIONAL	RFCxxxx
	SV6ENC	OPTIONAL	RFCxxxx
	IRLM	OPTIONAL	RFCxxxx
	IATYP	MANDATORY	RFCxxxx
	ISADDR	MANDATORY	RFCxxxx
	XRLM	OPTIONAL	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XSADDR	MANDATORY	RFCxxxx
	PORTMN	MANDATORY	RFCxxxx
	PORTMX	MANDATORY	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
npool		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	POOLID	MANDATORY	RFCxxxx
	POOLLT	OPTIONAL	RFCxxxx
	POOLHT	OPTIONAL	RFCxxxx
----	----	----	----
ngamht		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	GAMCNT	MANDATORY	RFCxxxx
----	----	----	----

ngapmht		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	GAPMCNT	MANDATORY	RFCxxxx
----	----	----	----
nsapmht		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
	SAPMCNT	MANDATORY	RFCxxxx
----	----	----	----
ngaml		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
----	----	----	----
ngapml		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	OPTIONAL	RFCxxxx
	DSUBIX	OPTIONAL	RFCxxxx
	PSRLM	MANDATORY	RFCxxxx
	PATYP	OPTIONAL	RFCxxxx
	PSADDR	OPTIONAL	RFCxxxx
----	----	----	----
ngsl		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
----	----	----	----
nsapml		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	SSUBIX	MANDATORY	RFCxxxx
----	----	----	----
nfpkt		OPTIONAL	RFCxxxx
	NATINST	OPTIONAL	RFCxxxx
	PSRLM	MANDATORY	RFCxxxx
	PATYP	MANDATORY	RFCxxxx
	PSADDR	MANDATORY	RFCxxxx
	PDADDR	MANDATORY	RFCxxxx
	SSUBIX	OPTIONAL	RFCxxxx

Table 17: NAT-Related STRUCTURED-DATA Registrations

9. References

9.1. Normative References

[I-D.behave-NAT-MIB]
Perreault, S., Tsou, T., and S. Sivakumar, "Additional Managed Objects for Network Address Translators (NAT) (Work in progress)", September 2013.

[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.

[RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, June 2000.

[RFC4265] Schliesser, B. and T. Nadeau, "Definition of Textual Conventions for Virtual Private Network (VPN) Management", RFC 4265, November 2005.

[RFC4363] Levi, D. and D. Harrington, "Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering, and Virtual LAN Extensions", RFC 4363, January 2006.

[RFC4784] Carroll, C. and F. Quick, "Verizon Wireless Dynamic Mobile IP Key Update for cdma2000(R) Networks", RFC 4784, June 2007.

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

[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.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[US-ASCII] American National Standards Institute, , "Coded Character Set -- 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.

9.2. Informative References

[I-D.behave-ipfix-nat-logging] Sivakumar, S. and R. Penno, "IPFIX Information Elements for logging NAT Events (Work in progress)", August 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)", July 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)", July 2013.

[I-D.softwire-map] Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, "Mapping of Address and Port with Encapsulation (MAP) (Work in progress)", August 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)", July 2013.

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

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[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.

[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.

[RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

[RFC7040] Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public IPv4-over-IPv6 Access Network", RFC 7040, November 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
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: tina.tsou.zouting@huawei.com

T. Taylor (editor)
Huawei Technologies
Ottawa
Canada

Email: tom.taylor.stds@gmail.com