Behave Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 06, 2014

Z. Chen
China Telecom
C. Zhou
Huawei Technologies
T. Tsou
Huawei Technologies (USA)
T. Taylor, Ed.
Huawei Technologies
October 03, 2013

Syslog Format for NAT Logging
draft-ietf-behave-syslog-nat-logging-04

Abstract

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

Copyright Notice

Table of Contents

1. Introduction

Operators already need to record the addresses assigned to subscribers at any point in time, for operational and regulatory reasons. When operators introduce NAT devices which support address sharing (e.g., Carrier Grade NATs (CGNs)) into their network, additional information has to be logged. This document and

[I-D.behave-ipfix-nat-logging] are provided in order to standardize the events and parameters to be recorded, using SYSLOG [RFC5424] and IPFIX [RFC5101] respectively. The content proposed to be logged by the two documents is exactly the same, but as will be seen, the choice of which to use in a given scenario is an engineering issue.

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

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

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

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

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

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

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

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

2. Deployment Considerations

2.1. Static and Dynamic NATs

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

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

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

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

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

2.2. Realms and Address Pools

A realm defines the scope within which a specific set of addresses are unique. In general these will be IPv4 or IPv6 addresses, but not necessarily. A counter-example specifically addressed by this document is the case of Gateway-Initiated DS-Lite [RFC6674], where individual host sites are identified by context identifiers of

various types. See further discussion in Section 2.3 and Section 2.3.1.

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

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

- multiple internal realms with overlapping address spaces;

- an external IPv4 public realm; and/or

- an external IPv6 public realm.

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

1. An address mapping is created between the internal realm and an external realm chosen based on information in the triggering packet or administrative request.

2. Using that address mapping, a transport binding is created between specific transport endpoints (e.g., between specific port values) in the two realms for the protocol required by the session, and added to the Binding Information Base (BIB).

3. Setup of the session is completed by mapping the destination address and port (if necessary) into the selected external realm.

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

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

2.2.1. Address Pools

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

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

From this description, one can see that ports are a limited resource, subject to exhaustion at the pool level and, with "paired" behaviour, at the level of the individual address. Log events are defined in Section 3.2.1 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 [I-D.softwire-public-4over6], MAP-E [I-D.softwire-map], and Lightweight 4over6 [I-D.softwire-lw4over6] have seen experimental deployment and are in the process of being standardized at the time of writing of this document.

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

A DS-Lite Address Family Transition Router (AFTR) includes a large-scale session-stateful NAT44 processing potentially millions of sessions per second. The special character of AFTR operation over

that of a traditional NAT44 is that the source IPv4 addresses of the internal hosts may not be unique. As a consequence, the session tables need to include an alternative identifier associated with the subscriber host. For basic DS-Lite, this will be the IPv6 address used to encapsulate the packets outgoing from the host. See Section 6.6 of [RFC6333]. For gateway-initiated DS-Lite [RFC6674], an identifier associated with the incoming tunnel from the host is used instead.

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

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

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

2.3.1. IP Addresses and Generalized Internal Addresses

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

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

o For traditional NATs, the source IPv4 address (for NAT44) or IPv6 address (for NAT64) is sufficient.

o For the DS-Lite, Lightweight 4over6 or MAP-E transition methods, the subscriber site can be identified by the IPv6 tunnel endpoint prefix or address provisioned to that site.

o Gateway-initiated DS-Lite uses the combination of a (typically) 32-bit context identifier (CID) and a softwire identifier (SWID). Several different realizations of these identifiers are described in Section 6 of [RFC6674]. From the point of view of this document, the SWID is represented by a realm identifier, leaving the CID as the value of the generalized internal address itself.

2.4. The Port Control Protocol (PCP)

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

2.5. Logging At the Customer Edge

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

3. NAT-Related Events and Parameters

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

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

The listed parameters include an optional reporting NAT type in each case. The reporting NAT type is an operator configured or implementation- defined string indicating the type of the reporting NAT (e.g., NAT44, DS-Lite AFTR). The same device can offer different functions depending on the particular packets being processed. The reporting NAT type is meant to be a hint to aid in interpretation of the event report.

3.1. Events Relating To Allocation Of Resources To Hosts

3.1.1. NAT Session Creation and Deletion

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

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

The following specific events are defined:

- 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 transport binding. The destination IP address and port and possibly the trigger are unique to the session. If the destination IP and port do not require remapping into the external realm, the internal values are redundant and SHOULD be omitted from the report. So long as the underlying BIB entry exists, the internal values can in any event be retrieved from the natMappingTable in the NAT MIB [I-D.Behave-NAT-MIB] using the combination of protocol, external realm, external destination address, and external destination port as key.

- Reporting NAT type (OPTIONAL);

- Internal realm (MANDATORY);

- Generalized internal address (MANDATORY);

- Internal port or ICMP identifier (MANDATORY);

- External realm (MANDATORY);

o External IP address (MANDATORY);

o External port or ICMP identifier (MANDATORY);

o Protocol identifier (MANDATORY);

o Internal destination IP address (as given in outgoing packets) (OPTIONAL);

o Internal destination port or ICMP identifier (as given in outgoing packets) (OPTIONAL);

o External destination IP address (as given in outgoing packets) (MANDATORY). It is unnecessary to specify the address type in the detailed encoding of this value, since the type will be the same as that of the external address parameter.

o External destination port or ICMP identifier (as given in outgoing packets) (MANDATORY);

o Trigger for session creation or deletion (OPTIONAL):

  * outgoing packet received;

  * incoming packet received;

  * administrative action (e.g., via the Port Control Protocol [RFC6887]); or

  * deletion of the underlying BIB entry.

3.1.1.1. **Destination Logging**

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

One possibility is that the implementation provides the operator with the ability to log destinations only for particular subscribers or

particular mapped addresses on a special study basis. This facility could be used for trouble-shooting or malicious activity tracing in particular cases as required. If such a capability is provided, the implementation MUST report session creation and deletion events for sessions matching the specified criteria, but MUST NOT report these events for other sessions.

3.1.2. Binding Information Base Entry Creation and Deletion

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

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

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

The following specific events are defined:

- o BIB entry creation
- o BIB entry deletion

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

- o Reporting NAT type (OPTIONAL);
- o Internal realm (MANDATORY);
- o Generalized internal address (MANDATORY);

o Internal port or ICMP identifier (MANDATORY);

o External realm (MANDATORY);

o External address (MANDATORY);

o External port or ICMP identifier (MANDATORY);

o Protocol identifier (MANDATORY);

o Trigger for transport binding creation or deletion (OPTIONAL):

   * outgoing packet received;

   * incoming packet received;

   * administrative action (e.g., via the Port Control Protocol [RFC6887]); or

   * deletion of the underlying address mapping.

3.1.3. Address Mapping Creation and Deletion Events

Two specific events are provided:

o Address mapping creation;

o Address mapping deletion.

Address mapping is discussed in detail in Section 2.2.

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

The address mapping events take the following specific parameters:

o Reporting NAT type (OPTIONAL);

o Internal realm (MANDATORY);

o Generalized internal address (MANDATORY);

o External realm (MANDATORY);

- o External IP address (MANDATORY).

- o Trigger for address mapping creation or deletion (OPTIONAL):

  - * outgoing packet;

  - * administrative action (e.g., via the Port Control Protocol [RFC6887]); or

  - * autonomous action of the NAT.

3.1.4. Port Set Allocation and Deallocation

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

The following specific events are defined:

- o Port set allocation;

- o Port set deallocation.

The parameters for these events are:

- o Reporting NAT type (OPTIONAL);

- o Internal realm (MANDATORY);

- o Generalized internal address (MANDATORY);

- o External realm (MANDATORY);

- o External IP address (MANDATORY);

- o A set of ports available for transport binding, newly allocated to or deallocated from the given address mapping. The representation of a port set is described in the next paragraph (MANDATORY).

- o Trigger for port set 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.

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

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

- o Ending port number, the highest port number in the entire port set (MANDATORY);

- o Range length, the number of port values in each range (OPTIONAL);

- o Range step, the difference between the first port number in one range and the first port number in the immediately preceding range of the port set (OPTIONAL).

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

Examples:

1. Two ranges, 1024-1535 and 2048-2559 are allocated. Each range consists of 512 consecutive port numbers. The parameter values to represent this allocation are:

   * starting port = 1024

   * ending port = 2559

   * range length = 512

   * range step = 1024.

2. Strictly for purposes of illustration, assume a sequence of 512 even-numbered ports is allocated, beginning at 1024, then 1026, ending at 2046. The parameter values to represent this allocation are:

   * starting port = 1024

   * ending port = 2046

   * range length = 1

   * range step = 2.

3. A single range of ports is allocated, running consecutively from 1024 to 2046. The parameter values to represent this allocation are:

   * starting port = 1024

   * ending port = 2046.

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

3.2. Events Directed Toward Operations and Maintenance

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

Two specific events provide reports on address pool utilization:

o High-water-mark threshold reached or exceeded;

o Low-water-mark threshold reached or under-shot.

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

Address pools are discussed in Section 2.2.1. The natPoolTable object in the NAT MIB [I-D.Behave-NAT-MIB] provides access to parameters describing the utilization level of address-port combinations within a given pool. Since a new transport mapping cannot be allocated unless a mappable address and a free port on that address are available, it is important to know when the available set of address-port combinations within a given pool is nearing exhaustion. Hence the natPoolTable contains a high-water-mark threshold settable by the operator. An address pool high-water-mark event report is generated when a new mapping into the pool is 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 wwhen 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 [RFC4787], 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 [I-D.Behave-NAT-MIB] dropped packet counters:

o globally, natResourceErrors in the natCounters table;

o per protocol, natProtocolResourceErrors in natProtocolTable;

o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

Packets triggering failed attempts to establish new connections due to port exhaustion are counted in the following NAT MIB [I-D.Behave-NAT-MIB] dropped packet counters:

o globally, natOutOfPortErrors in the natCounters table;

o per protocol, natProtocolOutOfPortErrors in natProtocolTable;

o per subscriber, natSubscriberOutOfPortErrors in natSubscribersTable.

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

o Reporting NAT type (OPTIONAL);

o Pool identifier, equal to the value of the natPoolIndex object presented in the natPoolTable in the MIB (MANDATORY).

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:

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

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

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

- o Reporting NAT type (OPTIONAL);

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

3.2.3. 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:

- o globally, natResourceErrors in the natCounters table;

- o per protocol, natProtocolResourceErrors in natProtocolTable;

- o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event are:

- o Reporting NAT type (OPTIONAL);

- o Trigger for address mapping creation (MANDATORY):

  - * outgoing packet;

   * administrative action (e.g., via the Port Control Protocol [RFC6887]).

3.2.4. Global BIB Entry High-Water-Mark Threshold Event

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

o BIB entry high-water-mark threshold exceeded.

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

The NAT MIB displays cumulative counts of mappings created in and removed from the BIB in the natCounters table. When the difference between these two counters 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:

o Reporting NAT type (OPTIONAL);

o Current number of active mappings, equal to the difference between the natMappingCreations and natMappingRemovals counters displayed in the natCounters table in the NAT MIB (MANDATORY).

3.2.5. Global BIB Entry Limit Exceeded

The global BIB entry limit exceeded event is reported when a new transport binding (i.e., BIB entry creation) is requested but the total number of transport bindings would exceed an administrative limit if it were added. The limit is given by object natLimitMappings in the natLimits table of the NAT MIB. MIB counters giving number of packets dropped due to resource limitations including this one are:

o globally, natResourceErrors in the natCounters table;

o per protocol, natProtocolResourceErrors in natProtocolTable;

o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event are:

o Reporting NAT type (OPTIONAL);

o Trigger for BIB entry creation (MANDATORY):

   * incoming packet;

   * outgoing packet;

   * administrative action (e.g., via the Port Control Protocol [RFC6887]).

3.2.6. Subscriber-Specific BIB Entry Threshold Event

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

o Subscriber-specific BIB entry high-water-mark threshold exceeded.

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

The NAT MIB displays cumulative counts of BIB entries 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 BIB entry high-water-mark threshold event is reported.

The specific parameters provided by this event report are:

o Reporting NAT type (OPTIONAL);

o Internal realm of the subscriber (MANDATORY);

o Generalized internal address of the subscriber (MANDATORY);

o Current number of active BIB entries for this subscriber, equal to the difference between the natSubscriberMappingCreations and natSubscriberMappingRemovals counters displayed in the natSubscribersTable table in the NAT MIB (MANDATORY).

3.2.7. 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 hosts 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:

o globally, natResourceErrors in the natCounters table;

o per protocol, natProtocolResourceErrors in natProtocolTable;

o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event are:

o Reporting NAT type (OPTIONAL);

o Trigger for mapping creation (MANDATORY):

  * outgoing packet;

  * administrative action (e.g., via the Port Control Protocol [RFC6887]).

3.2.8. Subscriber-Specific Limit On Number of BIB Entries Exceeded

The subscriber-specific limit on number of BIB entries exceeded event is reported when a new BIB entry is requested, but the total number of BIB entries for that subscriber would exceed an administrative limit if it were added. The limit is given by object natSubscriberLimitMappings in natSubscribersTable in the NAT MIB. MIB counters giving number of packets dropped due to resource limitations including this one are:

o globally, natResourceErrors in the natCounters table;

o per protocol, natProtocolResourceErrors in natProtocolTable;

o per subscriber, natSubscriberResourceErrors in natSubscribersTable.

The parameters for this event are:

o Reporting NAT type (OPTIONAL);

o Internal realm of the subscriber (MANDATORY);

o Generalized internal address of the subscriber (MANDATORY);

o Trigger for BIB entry creation (MANDATORY):

  * incoming packet;

  * outgoing packet;

   * administrative action (e.g., via the Port Control Protocol [RFC6887]).

3.2.9. Quota Exceeded

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

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

o globally, in counter natQuotaDrops in the natCounters table;

o by protocol, in the natProtocolQuotaDrops counter in the natProtocolTable;

o per subscriber, in counter natSubscriberQuotaDrops in the natSubscribersTable.

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

o Reporting NAT type (OPTIONAL);

o Quota identifier (MANDATORY);

o Internal realm (OPTIONAL);

o Generalized internal address (OPTIONAL);

o Source realm for triggering packet (OPTIONAL);

o Source IP address (OPTIONAL);

- o Source port or ICMP identifier (OPTIONAL);
- o Destination IP address (OPTIONAL);
- o Destination port (OPTIONAL);
- o Protocol (OPTIONAL);
- o Trigger for quota violation (OPTIONAL)
  - * packet received at the NAT;
  - * administrative action (e.g., via the Port Control Protocol [RFC6887]).

In the special case where the quota addresses bulk port allocation, the parameters listed above MUST be interpreted and populated as follows, so as to capture the address mapping to which the ports would have been allocated:

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

3.2.10. Global Limit On Number Of Fragments Pending Reassembly Exceeded

The global limit on number of fragments pending reassembly exceeded event is reported when a new fragment is received and the number of fragments currently awaiting reassembly is already equal to an administrative limit. That limit is given by the natLimitFragments object in the natLimits table. This event MUST NOT be reported unless the NAT supports the "receive fragments out of order" behavior [RFC4787]. MIB counters giving number of packets dropped due to resource limitations including this one are:

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

The parameters for this event provide the contents of the IP header of the received fragment that triggered it. If the source realm is

internal and the generalized internal address is available, it MUST also be included.

o Reporting NAT type (OPTIONAL);

o Source realm of the packet (MANDATORY);

o Source IP address (MANDATORY);

o Destination IP address (MANDATORY);

o Generalized internal address of the source (OPTIONAL).

4. SYSLOG Applicability

The primary advantage of SYSLOG is the human readability and searchability of its contents. In addition, it has built-in priority and 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 5.3.1.1.1, 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 BIB entry 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, BIB entries,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, limit, and quota 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. 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 "NATMTC" MUST be used for the event types defined in Section 5.3.2.

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 session creation	NAT	SADD	6 info	nsess
NAT session deletion	NAT	SDEL	6 info	nsess
BIB entry creation	NAT	BADD	6 info	nbib
BIB entry deletion	NAT	BDEL	6 info	nbib
Address mapping creation	NAT	AMADD	6 info	namap
Address mapping deletion	NAT	AMDEL	6 info	namap
Port set allocation	NAT	PTADD	6 info	npset
Port set deallocation	NAT	PTDEL	6 info	npset
Address pool high threshold	NATMTC	POOLHT	4 warning	npool
Address pool low threshold	NATMTC	POOLLT	6 info	npool
Global address map high threshold	NATMTC	GAMHT	4 warning	ngamht
Global address map limit	NATMTC	GAMLIM	3 error	ngaml
Global BIB entry high threshold	NATMTC	GBHT	4 warning	ngbht
Global BIB entry limit	NATMTC	GBLIM	3 error	ngbl
Subscriber-specific BIB entry high threshold	NATMTC	SBHT	5 notice	nsbht
Global active subscriber limit	NATMTC	GSLIM	3 error	ngsl
Subscriber-specific BIB entry limit	NATMTC	SBLIM	5 notice	nsbl
Quota exceeded	NATMTC	QUOTA	3-5 depending	nqpkt
Pending fragment limit	NATMTC	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, and the PARAM-NAMES and brief descriptions are listed in Table 2. They are then described more fully in the same order in succeeding sub-sections.

PARAM-NAME	Description
GAMCNT	Current global number of address mappings
GBCNT	Current global number of BIB entries
GIATYP	Generalized internal address type
GIAVAL	Generalized internal address/prefix value
IDATYP	Internal destination IP address type
IDAVAL	Internal destination IP address value
IDPNUM	Internal destination port or ICMP identifier value
IRLM	Internal realm
IPNUM	Internal port or ICMP identifier value (in BIB entry)
NTYP	Reporting NAT type
PDAVAL	Packet destination IP address value
PDPNUM	Packet destination port or ICMP identifier value
POOLID	Address pool identifier
PROTO	Protocol identifier
PSRLM	Packet source realm
PSATYP	Packet source IP address type
PSAVAL	Packet source IP address value
PSPNUM	Packet source port or ICMP identifier value
PTENUM	Port set ending number
PTSNUM	Port set starting number
QID	Quota identifier
RGLEN	Number of port values per range
RGSTEP	Difference between first values of successive port ranges
SBCNT	Current subscriber-specific number of active BIB entries
TRIG	Trigger for event
XATYP	External IP address type (in address mapping etc.)
XAVAL	External IP address value (in address mapping etc.)
XDAVAL	External destination IP address value (in session entry)
XDPNUM	External destination port or ICMP identifier value (in session entry)
XPNUM	External port or ICMP identifier value (in BIB entry)
XRLM	External realm (in address mapping etc.)

Table 2: Parameters Used In NAT-Related Log Reports, By PARAM-NAME

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.

5.2.2. PARAM-NAME GAMCNT: Current global number of address mappings

PARAM-VALUE: decimal number presented without leading zeroes.

Used with event types (MSGIDs): GAMHT.

5.2.3. PARAM-NAME GBCNT: Current global number of BIB entries

PARAM-VALUE: decimal number presented without leading zeroes.

Used with event types (MSGIDs): GBHT.

5.2.4. PARAM-NAME GIATYP: Generalized internal address type

PARAM-VALUE: enumeration giving the type of the generalized address. Possible values:

"IPv4": IPv4 address or prefix;

"IPv6": IPv6 address or prefix;

"GRE": Gateway-initiated DS-Lite [RFC6674] Context Identifier (CID) configured as a GRE key.

"MPLS": Gateway-initiated DS-Lite [RFC6674] Context Identifier (CID) configured as an MPLS label.

"FL": Gateway-initiated DS-Lite [RFC6674] Context Identifier (CID) configured as an IPv6 Flow Label.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL, SBHT, SBLIM, QUOTA, FRAG.

5.2.5. PARAM-NAME GIAVAL: Generalized internal address/prefix value

PARAM-VAL: If the value of GIATYP is "IPv4" or IPv6", the content of the GIAVAL parameter MUST be presented as an IPv4 or IPv6 address or prefix respectively as specified in Section 5.2.1. For all other types, the address MUST be presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL, SBHT, SBLIM, QUOTA, FRAG.

5.2.6. PARAM-NAME IDATYP: Internal destination IP address type

PARAM-VAL: IP address type. Possible values:

"IPv4": IPv4 address;

"IPv6": IPv6 address.

Used with event types (MSGIDs): SADD, SDEL.

5.2.7. PARAM-NAME IDAVAL: Internal destination IP address value

PARAM-VAL: IPv4 or IPv6 address, presented as specified in Section 5.2.1.

Used with event types (MSGIDs): SADD, SDEL.

5.2.8. PARAM-NAME IDPNUM: Internal destination port or ICMP identifier value

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): SADD, SDEL.

5.2.9. PARAM-NAME IRLM: Internal realm

PARAM-VAL: administratively-provided string of text.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL, SBHT, SBLIM, QUOTA.

5.2.10. PARAM-NAME IPNUM: Internal port or ICMP identifier value (in BIB entry)

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL.

5.2.11. PARAM-NAME NTYP: Reporting NAT type

PARAM-VAL: implementation- or operator-provided string of text.

Used with event types (MSGIDs): all.

5.2.12. PARAM-NAME PDAVAL: Packet destination IP address value

PARAM-VAL: IPv4 or IPv6 address, presented as specified in Section 5.2.1.

Used with event types (MSGIDs): QUOTA, FRAG.

5.2.13. PARAM-NAME PDPNUM: Packet destination port or ICMP identifier value

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): QUOTA.

5.2.14. PARAM-NAME POOLID: Address pool identifier

PARAM-VAL: 32-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): POOLHT, POOLLT.

5.2.15. PARAM-NAME PROTO: Protocol identifier

PARAM-VAL: A transport protocol number, from the "protocol-numbers" IANA registry, presented as a decimal number between 0 and 255 without leading zeroes.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, QUOTA.

5.2.16. PARAM-NAME PSRLM: Packet source realm

PARAM-VAL: administratively-provided string of text.

Used with event types (MSGIDs): QUOTA, FRAG.

5.2.17. PARAM-NAME PSATYP: Packet source IP address type

PARAM-VAL: IP address type. Possible values:

"IPv4": IPv4 address;

"IPv6": IPv6 address.

Used with event types (MSGIDs): QUOTA, FRAG.

5.2.18. PARAM-NAME PSAVAL: Packet source IP address value

PARAM-VAL: IPv4 or IPv6 address, presented as specified in Section 5.2.1.

Used with event types (MSGIDs): QUOTA, FRAG.

5.2.19. PARAM-NAME PSPNUM: Packet source port or ICMP identifier value

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): QUOTA.

5.2.20. PARAM-NAME PTENUM: Port set ending number

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): PTADD, PTDEL.

5.2.21. PARAM-NAME PTSNUM: Port set starting number

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): PTADD, PTDEL.

5.2.22. PARAM-NAME QID: Quota identifier

PARAM-VAL: 32-bit value presented as a decimal number without leading zeroes. [Note - ed.: have to confirm if and when MIB quota table is specified.]

Used with event types (MSGIDs): QUOTA.

5.2.23. PARAM-NAME RGLEN: Number of port values per range

PARAM-VAL: positive value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): PTADD, PTDEL.

5.2.24. PARAM-NAME RGSTEP: Difference between first values of successive port ranges

PARAM-VAL: up to 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): PTADD, PTDEL.

5.2.25. PARAM-NAME SBCNT: Current subscriber-specific number of active BIB entries

PARAM-VAL: value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): SBHT.

5.2.26. PARAM-NAME TRIG: Trigger for event

PARAM-VAL: enumeration. Possible values:

"OPKT": outgoing packet received at NAT.

"IPKT": incoming packet received at NAT.

"ADMIN": administrative action.

"BDEL": deletion of the underlying BIB entry.

"AMDEL": deletion of the underlying address mapping.

"AUTO": autonomous action of the NAT.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL, GAMLIM, GBLIM, GSLIM, SBLIM, QUOTA. Note that no event type supports all of the values listed above. The set of supported values is listed for each using event type in Section 5.3.

5.2.27. PARAM-NAME XATYP: External IP address type (in address mapping)

PARAM-VAL: IP address type. Possible values:

"IPv4": IPv4 address;

"IPv6": IPv6 address.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL.

5.2.28. PARAM-NAME XAVAL: External IP address value (in address mapping)

PARAM-VAL: IPv4 or IPv6 address, presented as specified in Section 5.2.1.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL.

5.2.29. PARAM-NAME XDAVAL: External destination IP address value

PARAM-VAL: IPv4 or IPv6 address, presented as specified in Section 5.2.1. Note that the type of address is given by XATYP, which will also be present in the event report.

Used with event types (MSGIDs): SADD, SDEL.

5.2.30. PARAM-NAME XDPNUM: External destination port or ICMP identifier value

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): SADD, SDEL.

5.2.31. PARAM-NAME XPNUM: External port or ICMP identifier value (in BIB entry)

PARAM-VAL: 16-bit value presented as a decimal number without leading zeroes.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL.

5.2.32. PARAM-NAME XRLM: External realm (in address mapping)

PARAM-VAL: administratively-provided string of text.

Used with event types (MSGIDs): SADD, SDEL, BADD, BDEL, AMADD, AMDEL, PTADD, PTDEL.

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. 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 Session Creation and Deletion

As shown in Table 1:

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

o 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 3. The requirements for these contents are derived from the description in Section 3.1.1.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
IRLM	Section 5.2.9	MANDATORY
GIATYP	Section 5.2.4	MANDATORY
GIAVAL	Section 5.2.5	MANDATORY
IPNUM	Section 5.2.10	MANDATORY
XRLM	Section 5.2.32	MANDATORY
XATYP	Section 5.2.27	MANDATORY
XAVAL	Section 5.2.28	MANDATORY
XPNUM	Section 5.2.31	MANDATORY
PROTO	Section 5.2.15	MANDATORY
IDATYP	Section 5.2.6	OPTIONAL
IDAVAL	Section 5.2.7	OPTIONAL
IDPNUM	Section 5.2.8	OPTIONAL
XDAVAL	Section 5.2.29	MANDATORY
XDPNUM	Section 5.2.30	MANDATORY

TRIG	Section 5.2.26	OPTIONAL

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

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", "BDEL", or "AUTO".

5.3.1.1.1. Example

This example is for a DS-Lite AFTR, hence the external addresses will be IPv4, as will the internal destination address. It is assumed that remapping of the destination address is unnecessary, so the internal values of that address are omitted. The generalized internal address is the IPv6 /56 prefix assigned to the site. Both the NTYP and TRIG optional parameters are present. The PRI value at the beginning of the log assumes a local use Facility value of 17 and Severity value 6. Note that the log could also include other SD-ELEMENTs (e.g., timeQuality).

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

```
<142>1 2013-05-07T22:14:15.03487Z record.example.net NAT 5063
SADD [nsess NTYP="AFTR" IRLM="MonteCristo-089" GIATYP="IPv6"
GIAVAL="2001:db8:a5e6:3900::/56" IPNUM="49178" XRLM="EXTv4"
XATYP="IPv4" XAVAL="198.51.100.127" XPNUM="6803"
PROTO="6" XDAVAL="192.0.2.57" TRIG="IPKT"]
```

Character count: about 270. Adding the internal destination address type, address value and port would bring this to 310.

5.3.1.2. BIB Entry Creation and Deletion

As shown in Table 1:

o NAT BIB entry creation event is indicated by MSGID set to "BADD";

o NAT BIB entry deletion event is indicated by MSGID set to "BDEL".

For both events, the associated SD-ELEMENT is tagged by SD-ID "nbib". The contents of the nbib SD-ELEMENT are shown in Table 4. The requirements for these contents are derived from the description in Section 3.1.2. The differences from the nsess SD-ELEMENT are the omission of the XDAVAL (external destination address) field in all cases and potentially the IDATYP and IDAVAL (internal destination address type and value) fields if mapping is required.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
IRLM	Section 5.2.9	MANDATORY
GIATYP	Section 5.2.4	MANDATORY
GIAVAL	Section 5.2.5	MANDATORY
IPNUM	Section 5.2.10	MANDATORY
XRLM	Section 5.2.32	MANDATORY
XATYP	Section 5.2.27	MANDATORY
XAVAL	Section 5.2.28	MANDATORY
XPNUM	Section 5.2.31	MANDATORY
PROTO	Section 5.2.15	MANDATORY
TRIG	Section 5.2.26	OPTIONAL

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

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

Using the same assumptions as in Section 5.3.1.1.1, the corresponding BIB entry creation report would look like this:

```
<142>1 2013-05-07T22:14:15.03487Z record.example.net NAT 5063
BADD [nbib NTYP="AFTR" IRLM="MonteCristo-089" GIATYP="IPv6"
GIAVAL="2001:db8:a5e6:3900::/56" IPNUM="49178" XRLM="EXTv4"
XATYP="IPv4" XAVAL="198.51.100.127" XPNUM="6803"
PROTO="6" TRIG="IPKT"]
```

Character count is about 255.

5.3.1.3. Address Mapping Creation and Deletion

As shown in Table 1:

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

o 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 5. The requirements for these contents are derived from the description in Section 3.1.3. The differences from the nbib SD-ELEMENT are the omission of the IPNUM, XPNUM, and PROTO (port number and protocol) fields.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
IRLM	Section 5.2.9	MANDATORY
GIATYP	Section 5.2.4	MANDATORY
GIAVAL	Section 5.2.5	MANDATORY
XRLM	Section 5.2.32	MANDATORY
XATYP	Section 5.2.27	MANDATORY
XAVAL	Section 5.2.28	MANDATORY
TRIG	Section 5.2.26	OPTIONAL

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

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".

Again using the same assumptions as in Section 5.3.1.1.1, but assuming the address mapping was created earlier, the corresponding address mapping entry creation report would look like this:

```
<142>1 2013-05-07T22:14:12.95628Z record.example.net NAT 5063
AMADD [namap NTYP="AFTR" IRLM="MonteCristo-089" GIATYP="IPv6"
GIAVAL="2001:db8:a5e6:3900::/56" XRLM="EXTv4"
XATYP="IPv4" XAVAL="198.51.100.127" TRIG="OPKT"]
```

Character count is about 225.

5.3.1.4. Port Set Allocation and Deallocation

As shown in Table 1:

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

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

For both events, the associated SD-ELEMENT is tagged by SD-ID "npset". The contents of the npset SD-ELEMENT are shown in Table 6. The requirements for these contents are derived from the description in Section 3.1.4.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
IRLM	Section 5.2.9	MANDATORY
GIATYP	Section 5.2.4	MANDATORY
GIAVAL	Section 5.2.5	MANDATORY
XRLM	Section 5.2.32	MANDATORY
XATYP	Section 5.2.27	MANDATORY
XAVAL	Section 5.2.28	MANDATORY
PTSNUM	Section 5.2.21	MANDATORY
PTSNUM	Section 5.2.21	MANDATORY
RGLEN	Section 5.2.23	OPTIONAL
RGSTEP	Section 5.2.24	OPTIONAL
TRIG	Section 5.2.26	OPTIONAL

Table 6: Contents Of the SD-ELEMENT Section For Logging the Port Set Allocation and Deallocation Events

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 the first example in Section 3.1.4, where two ranges, 1024-1535 and 2048-2559 are allocated to the address mapping on which the example in Section 5.3.1.3 is based. The corresponding port set allocation report would look like this:

```
<142>1 2013-08-15T09:14:38.12229Z record.example.net NAT 5063
PTADD [npset NTYP="AFTR" IRLM="MonteCristo-089"
GIATYP="IPv6" GIAVAL="2001:db8:a5e6:3900::/56" XRLM="EXTv4"
XATYP="IPv4" XAVAL="198.51.100.127" PTSNUM="1024" PTENUM="2559"
RGLEN="512" RGSTEP="1024" TRIG="IPKT"]
```

Character count is about 270.

5.3.2. Events Directed Toward Operations and Maintenance

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

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

As shown in Table 1:

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

- o 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 7. The requirements for these contents are derived from the description in Section 3.2.1.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
POOLID	Section 5.2.14	MANDATORY

Table 7: Contents Of the SD-ELEMENT Section For Logging the Address Pool High- and Low-Water-Mark Threshold Events

Example, assuming a local-use Facility value of 16 and a Severity level of 4 (warning) to calculate the PRI value at the beginning:

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

Character count is about 100.

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

As shown in Table 1:

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

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

The contents of the ngamht SD-ELEMENT are shown in Table 8. The requirements for these contents are derived from the description in Section 3.2.2.

PARAM-NAME	Description	Requirement

NTYP	Section 5.2.11	OPTIONAL
GAMCNT	Section 5.2.2	MANDATORY

Table 8: Contents Of the SD-ELEMENT Section For Logging the Global Address Map High-Water-Mark Threshold Event

Example, assuming a local-use Facility value of 16 and a Severity level of 4 (warning) to calculate the PRI value at the beginning. Suppose the threshold was set to 690000, so it has already been 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 NATMTC 5025
GAMHT [ngamht NTYP="AFTR" GAMCNT="690015"]
```

Character count is about 100.

5.3.2.3. 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 9. The requirements for these contents are derived from the description in Section 3.2.3.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
TRIG	Section 5.2.26	MANDATORY

Table 9: Contents Of the SD-ELEMENT Section For Logging the Global Address Map Limit Exceeded Event

For the global address map limit exceeded event, TRIG can take on the values "OPKT" or "ADMIN".

Example, assuming a local-use Facility value of 16 and a Severity level of 3 (error) to calculate the PRI value at the beginning.

```
<131>1 2013-08-15T09:15:16.08716Z record.example.net NATMTC 5025
GAMLIM [ngaml NTYP="AFTR" TRIG="OPKT"]
```

Character count is about 100.

5.3.2.4. Global BIB Entry High-Water-Mark Threshold Event

As shown in Table 1:

- o Global BIB entry high-water-mark threshold event is indicated by MSGID set to "GBHT"; and
- o the associated SD-ELEMENT is tagged by SD-ID "ngbht".

The contents of the ngbht SD-ELEMENT are shown in Table 10. The requirements for these contents are derived from the description in Section 3.2.4.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
GBCNT	Section 5.2.3	MANDATORY

Table 10: Contents Of the SD-ELEMENT Section For Logging the Global BIB Entry High-Water-Mark Threshold Event

Example, assuming a local-use Facility value of 16 and a Severity level of 4 (warning) to calculate the PRI value at the beginning. Suppose the threshold was set to 2000000, so it has already been 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 NATMTC 5025
GBHT [ngbht NTYP="AFTR" GBCNT="2000023"]
```

Character count is about 100.

5.3.2.5. Global BIB Entry Limit Exceeded

As shown in Table 1:

- o Global BIB entry limit exceeded event is indicated by MSGID set to "GBLIM"; and
- o the associated SD-ELEMENT is tagged by SD-ID "ngbl".

The contents of the ngbl SD-ELEMENT are shown in Table 11. The requirements for these contents are derived from the description in Section 3.2.5.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
TRIG	Section 5.2.26	MANDATORY

Table 11: Contents Of the SD-ELEMENT Section For Logging the Global BIB Entry Limit Exceeded Event

For the global BIB entry limit exceeded event, TRIG can take on the values "OPKT", "IPKT", or "ADMIN".

Example, assuming a local-use Facility value of 16 and a Severity level of 3 (error) to calculate the PRI value at the beginning.

```
<131>1 2013-08-15T09:15:16.08Z record.example.net NATMTC 5025
GBLIM [ngbl NTYP="AFTR" TRIG="OPKT"]
```

Character count is about 100.

5.3.2.6. Subscriber-Specific BIB Entry High-Water-Mark Threshold Event

As shown in Table 1:

- o Subscriber-specific BIB entry high-water-mark threshold event is indicated by MSGID set to "SBHT"; and
- o the associated SD-ELEMENT is tagged by SD-ID "nsbht".

The contents of the nsbht SD-ELEMENT are shown in Table 12. The requirements for these contents are derived from the description in Section 3.2.6.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
IRLM	Section 5.2.9	MANDATORY
GIATYP	Section 5.2.4	MANDATORY
GIAVAL	Section 5.2.5	MANDATORY
SBCNT	Section 5.2.25	MANDATORY

Table 12: Contents Of the SD-ELEMENT Section For Logging the Subscriber-Specific BIB Entry High-Water-Mark Threshold Event

Example, assuming a local-use Facility value of 16 and a Severity level of 5 (notice) to calculate the PRI value at the beginning. Suppose the threshold was set to 1500 and the number of BIB entries 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 NATMTC 5025
SBHT [nsbht SBCNT="1501" IRLM="MonteCristo-089"
GIATYP="IPv6" GIAVAL="2001:db8:a5e6:3900::/56"]
```

Character count is about 155.

5.3.2.7. Global Limit On Number of Active Hosts Exceeded

As shown in Table 1:

- o 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 13. The requirements for these contents are derived from the description in Section 3.2.7.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
TRIG	Section 5.2.26	MANDATORY

Table 13: Contents Of the SD-ELEMENT Section For Logging the Global Active Host Limit Exceeded Event

For the global active host limit exceeded event, TRIG can take on the values "OPKT" or "ADMIN".

Example, assuming a local-use Facility value of 16 and a Severity level of 3 (error) to calculate the PRI value at the beginning.

```
<131>1 2013-08-15T09:15:16.08421Z record.example.net NATMTC 5025
GSLIM [ngsl NTYP="AFTR" TRIG="OPKT"]
```

Character count is about 95.

5.3.2.8. Subscriber-Specific Limit On Number of BIB Entries Exceeded

As shown in Table 1:

- o Subscriber-specific BIB entry limit exceeded event is indicated by MSGID set to "SBLIM"; and

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

The contents of the nsbl SD-ELEMENT are shown in Table 14. The requirements for these contents are derived from the description in Section 3.2.8.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
IRLM	Section 5.2.9	MANDATORY
GIATYP	Section 5.2.4	MANDATORY
GIAVAL	Section 5.2.5	MANDATORY
TRIG	Section 5.2.26	MANDATORY

Table 14: Contents Of the SD-ELEMENT Section For Logging the Subscriber-Specific BIB Entry Limit Exceeded Event

For the subscriber-specific BIB entry limit exceeded event, TRIG can take on the values "OPKT", "IPKT", or "ADMIN".

Example, assuming a local-use Facility value of 16 and a Severity level of 4 (warning) to calculate the PRI value at the beginning.

```
<132>1 2013-08-15T09:15:16.08528Z record.example.net NATMTC 5025
SBLIM [nsbl NTYP="AFTR" IRLM="MonteCristo-089"
GIATYP="IPv6" GIAVAL="2001:db8:a5e6:3900::/56" TRIG="OPKT"]
```

Character count is about 170.

5.3.2.9. Quota Exceeded

As shown in Table 1:

- o Quota exceeded event is indicated by MSGID set to "QUOTA"; and

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

The contents of the nqpkt SD-ELEMENT are shown in Table 15. The requirements for these contents are derived from the description in Section 3.2.9.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
QID	Section 5.2.22	MANDATORY
IRLM	Section 5.2.9	OPTIONAL
GIATYP	Section 5.2.4	OPTIONAL
GIAVAL	Section 5.2.5	OPTIONAL
PSRLM	Section 5.2.16	OPTIONAL
PSATYP	Section 5.2.17	OPTIONAL
PSAVAL	Section 5.2.18	OPTIONAL
PSPNUM	Section 5.2.19	OPTIONAL
PDAVAL	Section 5.2.18	OPTIONAL
PDPNUM	Section 5.2.19	OPTIONAL
PROTO	Section 5.2.15	OPTIONAL
TRIG	Section 5.2.26	OPTIONAL

Table 15: Contents Of the SD-ELEMENT Section For Logging the Quota Exceeded Event

For the quota exceeded event, TRIG can take on the values "OPKT", "IPKT", or "ADMIN".

First example, assuming a local-use Facility value of 16 and a Severity level of 4 (warning) to calculate the PRI value at the beginning. The quota was triggered by the arrival of a UDP/IPv4 packet from the exterior. An address mapping already exists, so that the generalized internal address corresponding to the packet destination is known and must be presented.

```
<132>1 2013-08-15T09:15:16.08Z record.example.net NATMTC 5025
QUOTA [nqpkt NTYP="AFTR" QID="21" IRLM="MonteCristo-089"
GIATYP="IPv6" GIAVAL="2001:db8:a5e6:3900::/56" PROTO="17"
PSRLM="EXTv4" PSATYP="IPv4" PSAVAL="203.0.113.26" PSPNUM="9803"
PDAVAL="198.51.100.127" PDPNUM="49853" TRIG="IPKT"]
```

Character count is about 290.

Second example, assuming a local-use Facility value of 16 and a Severity level of 5 (notice) to calculate the PRI value at the beginning. The quota was triggered by a PCP request based on [I-D.pcp-port-set] to allocate more ports to an existing address mapping. Since the address mapping already exists, the generalized

internal address corresponding to the request is known and must be presented.

```
<133>1 2013-08-15T09:15:16.08Z record.example.net NATMTC 5025
QUOTA [nqpkt NTYP="AFTR" QID="48" IRLM="MonteCristo-089"
GIATYP="IPv6" GIAVAL="2001:db8:a5e6:3900::/56"
PSRLM="EXTv4" PSATYP="IPv4" PSAVAL="198.51.100.127" TRIG="ADMIN"]
```

Character count is about 230.

5.3.2.10. Pending Fragment Limit Exceeded

As shown in Table 1:

- o Pending fragment limit exceeded event is indicated by MSGID set to "FRAG"; and
- o 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 3.2.10.

PARAM-NAME	Description	Requirement
NTYP	Section 5.2.11	OPTIONAL
PSRLM	Section 5.2.16	MANDATORY
PSATYP	Section 5.2.17	MANDATORY
PSAVAL	Section 5.2.18	MANDATORY
PDAVAL	Section 5.2.12	MANDATORY
GIATYP	Section 5.2.4	OPTIONAL
GIAVAL	Section 5.2.5	OPTIONAL

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

Example, assuming a local-use Facility value of 16 and a Severity level of 4 (warning) to calculate the PRI value at the beginning. 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 NATMTC 5025
FRAG [nfpkt NTYP="AFTR" PSRLM="MonteCristo-089"
PSATYP="IPv4" PSAVAL="192.0.0.1" PDAVAL="203.0.113.26"
GIATYP="IPv6" GIAVAL="2001:db8:a5e6:3900::/56"]
```

Character count is about 210.

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:

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

- o 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.

- o 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="NATMTC", MSGID="POOLHT") is obviously more urgent than the low-water-mark threshold event (APP-NAME="NATMTC", 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="NATMTC"). 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 Suppress Event Reports

The event report types specified with APP-NAME="NATMTC" all relate to limits or thresholds. By their nature, events of this sort will come in bursts. The limit or threshold 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 when to suppress individual event report types, all implementations MUST allow the operator to indicate through configuration that a given event report type is to be completely suppressed (i.e., disabled). This is particularly required to disable destination logging when that is not required (see Section 3.1.1.1). It is also required when the operator prefers to receive particular event notifications via SNMP rather than SYSLOG.

The ability to suppress 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 suppressed messages 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 "NATMTC" 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:

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

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

o 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. [Ed. note -- if we can get experience or simulation results we may be able to add ballpark figures.]

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.

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 5.3.2 are less sensitive, but the subscriber-specific threshold and limit events reveal internal realm and generalized internal address information which might be of interest to outside attackers. The quota event and the fragmentation limit event also provide 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.

Structured Data ID	Structured Data Parameter	Required or Optional	Reference
nsess		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	IRLM	MANDATORY	RFCxxxx
	GIATYP	MANDATORY	RFCxxxx
	GIAVAL	MANDATORY	RFCxxxx
	IPNUM	MANDATORY	RFCxxxx

	XRLM	MANDATORY	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XAVAL	MANDATORY	RFCxxxx
	XPNUM	MANDATORY	RFCxxxx
	PROTO	MANDATORY	RFCxxxx
	IDATYP	OPTIONAL	RFCxxxx
	IDAVAL	OPTIONAL	RFCxxxx
	IDPNUM	OPTIONAL	RFCxxxx
	XDAVAL	MANDATORY	RFCxxxx
	XDPNUM	MANDATORY	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
nbib		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	IRLM	MANDATORY	RFCxxxx
	GIATYP	MANDATORY	RFCxxxx
	GIAVAL	MANDATORY	RFCxxxx
	IPNUM	MANDATORY	RFCxxxx
	XRLM	MANDATORY	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XAVAL	MANDATORY	RFCxxxx
	XPNUM	MANDATORY	RFCxxxx
	PROTO	MANDATORY	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
namap		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	IRLM	MANDATORY	RFCxxxx
	GIATYP	MANDATORY	RFCxxxx
	GIAVAL	MANDATORY	RFCxxxx
	XRLM	MANDATORY	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XAVAL	MANDATORY	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
npset		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	IRLM	MANDATORY	RFCxxxx
	GIATYP	MANDATORY	RFCxxxx
	GIAVAL	MANDATORY	RFCxxxx
	XRLM	MANDATORY	RFCxxxx
	XATYP	MANDATORY	RFCxxxx
	XAVAL	MANDATORY	RFCxxxx
	PTSNUM	MANDATORY	RFCxxxx
	PTSNUM	MANDATORY	RFCxxxx
	RGLEN	OPTIONAL	RFCxxxx
	RGSTEP	OPTIONAL	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx

----	----	----	----
npool		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	POOLID	MANDATORY	RFCxxxx
----	----	----	----
ngamht		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	GAMCNT	MANDATORY	RFCxxxx
----	----	----	----
ngaml		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	TRIG	MANDATORY	RFCxxxx
----	----	----	----
ngbht		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	GBCNT	MANDATORY	RFCxxxx
----	----	----	----
ngbl		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	TRIG	MANDATORY	RFCxxxx
----	----	----	----
nsbht		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	IRLM	MANDATORY	RFCxxxx
	GIATYP	MANDATORY	RFCxxxx
	GIAVAL	MANDATORY	RFCxxxx
	SBCNT	MANDATORY	RFCxxxx
----	----	----	----
ngsl		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	TRIG	MANDATORY	RFCxxxx
----	----	----	----
nsbl		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	IRLM	MANDATORY	RFCxxxx
	GIATYP	MANDATORY	RFCxxxx
	GIAVAL	MANDATORY	RFCxxxx
	TRIG	MANDATORY	RFCxxxx
----	----	----	----
nqpkt		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	QID	MANDATORY	RFCxxxx
	IRLM	OPTIONAL	RFCxxxx
	GIATYP	OPTIONAL	RFCxxxx
	GIAVAL	OPTIONAL	RFCxxxx
	PSRLM	OPTIONAL	RFCxxxx
	PSATYP	OPTIONAL	RFCxxxx
	PSAVAL	OPTIONAL	RFCxxxx

	PSPNUM	OPTIONAL	RFCxxxx
	PDAVAL	OPTIONAL	RFCxxxx
	PDPNUM	OPTIONAL	RFCxxxx
	PROTO	OPTIONAL	RFCxxxx
	TRIG	OPTIONAL	RFCxxxx
----	----	----	----
nfpkt		OPTIONAL	RFCxxxx
	NTYP	OPTIONAL	RFCxxxx
	PSRLM	MANDATORY	RFCxxxx
	PSATYP	MANDATORY	RFCxxxx
	PSAVAL	MANDATORY	RFCxxxx
	PDAVAL	MANDATORY	RFCxxxx
	GIATYP	OPTIONAL	RFCxxxx
	GIAVAL	OPTIONAL	RFCxxxx

Table 17: NAT-Related STRUCTURED-DATA Registrations

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.

[RFC2685] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", RFC 2685, September 1999.

[RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.

[RFC5425] Miao, F., Ma, Y., and J. Salowey, "Transport Layer Security (TLS) Transport Mapping for Syslog", RFC 5425, March 2009.

[RFC5848] Kelsey, J., Callas, J., and A. Clemm, "Signed Syslog Messages", RFC 5848, May 2010.

[RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

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

[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.softwire-public-4over6]
Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public IPv4 over IPv6 Access Network (Work in progress)", July 2013.

[I-D.softwire-unified-cpe]
Boucadair, M. and I. Farrer, "Unified IPv4-in-IPv6 Softwire CPE (Work in progress)", May 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.

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

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

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

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

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

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

[RFC6674] Brockners, F., Gundavelli, S., Speicher, S., and D. Ward, "Gateway-Initiated Dual-Stack Lite Deployment", RFC 6674, July 2012.

[RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

[RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, April 2013.

Authors' Addresses

Zhonghua Chen
China Telecom
P.R. China

Email: 18918588897@189.cn

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

Email: cathy.zhou@huawei.com

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

Phone: +1 408 330 4424
Email: tina.tsou.zouting@huawei.com

T. Taylor (editor)
Huawei Technologies
Ottawa
Canada

Email: tom.taylor.stds@gmail.com