Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 15, 2014

Z. Chen
China Telecom
C. Zhou
Huawei Technologies
T. Tsou
Huawei Technologies (USA)
T. Taylor
Huawei Technologies

July 14, 2013

Syslog Format for NAT Logging
draft-ietf-behave-syslog-nat-logging-02

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 term "Session" as it is defined in Section 2.3 of [RFC2663] and the term Binding Information Base (BIB) as it is defined in Section 2 of [RFC6146].

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).

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. 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 interior 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. This document does not include any requirements specific to the B4, since logs are not usually collected from customer equipment.

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.

The border router (BR), for either Lightweight 4over6 or MAP-E, is required to monitor port usage by outgoing IPv4 packets. If the ports used by a host fall outside its configured port set, the border router may return an ICMPv6 type 1, code 5 (source address failed ingress/egress policy) error message to the unified CPE. It is also possible for the same reason that the unified CPE receives incoming IPv4 packets with destination port numbers outside of its assigned range.

Out-of-range ports are a sign of misconfiguration or other problems, so it is reasonable for the BR to log such events (subject to the rate limiting). The log should capture the port set that the BR believes is configured on the unified CPE. For both Lightweight 4over6 and MAP-E, this is associated with an identifier, the 16-bit port set identifier (PSID).

2.3. 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.

2.4. 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. 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 device identifier and an optional reporting device type in each case. The reporting device identifier is potentially useful only if the HOSTNAME field in the log header identifies an off-board device rather than the NAT itself. The reporting device type identifies which of the reporting device types listed in Section 5.2.3 is reporting the event.

Reference will be made below to a subscriber site identifier. IOn practice, NATs use various means to distinguish customer endpoints, and this will be reflected in what they log. From a strictly theoretical point of view:

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

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

o Gateway-initiated DS-Lite uses the combination of a 32-bit context identifier (CID) and a softwire identifier (SWID). Several different realizations of these identifiers are described in Section 6 of [RFC6674].

3.1. NAT Session Creation and Deletion

NAT session creation and deletion events may be logged in a fully dynamic NAT when a binding from a subscriber site identifier and source port to an external address and port is recorded in or deleted from the session database. See Section 3 of [RFC3022] for more details about session creation and deletion.

The following specific events are defined:

o NAT Session Creation

o NAT Session Deletion

These take the same parameters for all types of NAT, aside from the variation in subscriber site identifier noted above:

- o reporting device type (OPTIONAL);
- o reporting device identifier (OPTIONAL);
- o Subscriber site identifier (MANDATORY);
- o Mapped external IPv4 address (MANDATORY);
- o Protocol identifier (MANDATORY for NAPT);
- o Internal port or ICMP identifier (MANDATORY for NAPT);
- o Mapped external port or ICMP identifier (MANDATORY for NAPT);
- o Address realm (internal or external) of the source of the packet triggering the creation of the session (OPTIONAL).

3.1.1. Destination Logging

The logging of destination address and port for outgoing packets is considered out of scope of this document, 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 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 short, destination logging will be a rarely-used procedure for which standardization seems unnecessary.

3.2. Address Binding Event

This event is recorded at a dynamic or hybrid NAT when a given subscriber site identifier has been bound to an external source address. An address binding occurs when the first packet in the first flow from the host in the internal realm is received at the NAT. It MAY occur under other circumstances (e.g., PCP request, or NAT policy permits assignment of a new external address due to port conflict). The event parameters are:

- o reporting device type (OPTIONAL);
- o reporting device identifier (OPTIONAL);
- o Subscriber site identifier (MANDATORY);

- Mapped external IPv4 address (MANDATORY).

3.3. Port Allocation Change

This event is recorded at a hybrid NAT whenever the set of ports allocated to a given address binding changes. When ports are allocated, the same ports are allocated for UDP and for TCP. The parameters for this event are:

- reporting device type (OPTIONAL);
- reporting device identifier (OPTIONAL);
- Subscriber site identifier (MANDATORY);
- Mapped external IPv4 address (MANDATORY);
- One or more contiguous port ranges specified by starting and ending port number (MANDATORY).

The log MUST indicate the cumulative set of ports allocated to the address binding (taking account both of allocations and deallocations) at the time the log was generated. An implementation MAY show each individual allocation as a separate range, or MAY consolidate adjacent ranges. For example, suppose the address binding is initially allocated the range 1024-1535. The log at the time of the initial allocation will contain that one range. Suppose now that an additional allocation is granted, consisting of the range 1536-2047. The log generated may contain the two ranges 1024-1535 and 1536-2047, or may contain the one consolidated range 1024-2047. Finally, suppose that ports 1536-1791 are deallocated. The resulting log will show the ranges 1024-1535 and 1792-2047 as the current allocation to the address binding.

3.4. NAT Address Exhaustion Event

This event will be generated when a NAT device runs out of global IPv4 addresses in a given pool of addresses. Typically, this event would mean that the NAT device will not be able to create any new translations until some addresses or ports are freed. This event takes the following parameters:

- reporting device type (OPTIONAL);
- reporting device identifier (OPTIONAL);
- address pool identifier (MANDATORY).

Implementations MUST provide the ability to limit the rate at which this log is generated, since the NAT may move back and forth between exhausted and almost-exhausted state many times during a particular busy episode.

3.5. Port Exhaustion Event

This event will be generated when a NAT device runs out of ports for a global IPv4 address. Port exhaustion shall be reported per protocol (UDP, TCP) individually. The event parameters are:

- reporting device type (OPTIONAL);
- reporting device identifier (OPTIONAL);
- Mapped external IPv4 address (MANDATORY);
- Protocol identifier (MANDATORY).

Implementations MUST provide the ability to limit the rate at which this log is generated, since the NAT may move back and forth between exhausted and almost-exhausted state many times during a particular busy episode.

3.6. Quota Exceeded Event

A "Quota Exceeded" event is reported when the NAT cannot allocate a new session because of an administratively imposed limit on the number of sessions allowed for a given subscriber or set of subscribers, for a given protocol or totalled over all protocols. The parameters of this event are:

- reporting device type (OPTIONAL);
- reporting device identifier (OPTIONAL);
- Site scope (MANDATORY);
- Protocol (MANDATORY);
- Subscriber site identifier (OPTIONAL);
- VLAN identifier or VPN Routing and Forwarding (VRF) identifier (OPTIONAL).

Site scope is either single site, multiple sites served by the same VLAN or VRF, or all sites served by the NAT. If the site scope is single site, then the subscriber site identifier MUST be present and

the VLAN or VRF identifier MUST be absent. If the site scope is multiple sites, then the reverse MUST be true. If the site scope is all sites, the subscriber site identifier, the VRF idenbtifier, and the VLAN identifier MUST NOT be present. Protocol scope is either a specific protocol (UDP, TCP, ICMP) or all protocols, meaning that the quota concerned applies to the total number of sessions supported by the NAT regardless of protocol.

Implementations MUST provide the ability to limit the rate at which this log is generated, since the NAT may move back and forth between over-quota and within-quota state many times during a particular busy episode.

3.7. Invalid Port Detected

As discussed in Section 2.2, this event may be reported at MAP-E or Lightweight 4over6 Border Router, either through receipt of ICMP error messages or by direct observation of incoming IPv4 packets.

> List discussion has pointed out that enabling ICMP on the customer edge device opens up the potential for denial of service attacks on the Border Router. Hence direct observation will be the more likely trigger for this event.

The event report takes the following parameters:

- reporting device identifier (OPTIONAL);

- Subscriber site identifier (MANDATORY);

- port set identifier for the subscriber site, as provisioned at the Border Router (MANDATORY).

3.8. Static NAT Configuration Event

*** Should we anticipate logging of the configuration (IPv6 prefix, IPv4 prefix/address, PSID) assigned to the Lightweight 4over6 or MAP-E CPE? ***

4. SYSLOG Applicability

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

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

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

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

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

5. SYSLOG Record Format For NAT Logging

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

5.1. SYSLOG HEADER Fields

Within the HEADER portion of the SYSLOG record, the priority (PRI) level is subject to local policy, but a default value of 8x is suggested, representing a Facility value of 10 (security/authorization) and a Severity level varying with the event type. The suggested value by event type is shown in Table 1. Depending on where the SYSLOG record is generated, the HOSTNAME field may identify the NAT or an offline logging device. In the latter case, it may be desirable to identify the NAT using the DevID field in the STRUCTURED-DATA section (see below). The value of the HOSTNAME field is subject to the preferences given in Section 6.2.4 of [RFC5424].

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

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

Table 1: Recommended MSGID Encodings and Default PRI Values for the Events Defined In Section 3

5.2. Parameter Encodings

This section describes how to encode the individual parameters that can appear in NAT-related logs. The parameters are taken from the event descriptions in Section 3, and are listed in Table 2. Formally, as will be seen in Table 10, a parameter used with more than one event is registered as multiple separate parameters, one for each event report in which it is used. However, there is no reason to change either the PARAM-NAME or the encoding of the PARAM-VALUE between different instances of the same parameter.

PARAM-NAME	Parameter
APoolId	Address pool identifier
DevID	reporting device identifier
DevTyp	reporting device type
PostS4	Mapped external IPv4 address
PostSPt	Mapped external port or ICMP identifier
PreSPt	Internal port or ICMP identifier
Proto	Protocol identifier
PScop	Protocol scope for quota
PSID	Port set identifier
PtRg	Range of consecutive port numbers
SiteID	Subscriber site identifier
SScop	Site scope for quota
TrigR	Address realm triggering the creation of the session
VLANid	VLAN identifier
VRFid	VPN routing and forwarding identifier

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

5.2.1. APoolId: Address Pool Identifier

PARAM-Value: decimal integer identifying a specific address pool at the reporting NAT.

5.2.2. DevID: Reporting Device Identifier

PARAM-VALUE: a UTF-8 string identifying the NAT or BR observing the event which this record reports. Needed only if the necessary identification is not provided by the HOSTNAME parameter in the log record header.

5.2.3. DevTyp: Reporting Device Type

PARAM-VALUE: one of the values provided in the IANA SYSLOG reporting device type registry established by this document. The initial values in that registry are:

44 NAT44 [RFC3022];

64 NAT64 [RFC6145] or [RFC6146];

AFTR DS-Lite AFTR [RFC6333];

BR Lightweight 4over6 or MAP-E border router.

This parameter is primarily additional information for the human reader of a log report, but could be used to provide a consistency check on the contents of a log. Instances where parameter usage depends on the reporting device type of the reporting NAT are noted in Section 5.3.

5.2.4. PostS4: Mapped External IPv4 Address

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

5.2.5. PostSPt: Mapped External Port or ICMP Identifier

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

5.2.6. PreSPt: Internal Port or ICMP Identifier

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

5.2.7. Proto: Protocol Identifier

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

5.2.8. PScop: Protocol Scope For Quota

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

5.2.9. PSID: Port Set Identifier

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

5.2.10. PtRg: Allocated Port Range

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

5.2.11. SiteID: Subscriber Site Identifier

A human-readable UTF-8 string identifying a specific host or CPE served by the reporting device. The type of identifier depends on the configuration of the reporting device, and is implementation and deployment-specific. See Section 3 for a discussion of the possible identifier types.

5.2.12. SScop: Site Scope For Quota

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

5.2.13. TrigR: Realm Triggering Session Creation

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

5.2.14. VLANid: VLAN Identifier

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

5.2.15. VRFid: VPN Routing and Forwarding Identifier

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

5.3. Encoding Of Complete Log Report For Each Event Type

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

5.3.1. NAT Session Creation and Deletion

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

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

PostS4	Section 5.2.4	MANDATORY
Proto	Section 5.2.7	MANDATORY
PreSPt	Section 5.2.6	MANDATORY
PostSPt	Section 5.2.5	MANDATORY
TrigR	Section 5.2.13	OPTIONAL

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

5.3.1.1. Examples

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

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

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

Character count: about 205.

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

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

The character count: about 165.

5.3.2. Address Binding Event

As shown in Table 1, the NAT address binding event is indicated by MSG-ID set to "AddrBind". The associated SD-ELEMENT is tagged by SD-ID "NATBind". The contents of the NATBind SD-ELEMENT are shown in Table 4. The requirements for these contents are derived from the description in Section 3.2.

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

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

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

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

Character count: about 125.

5.3.3. Port Allocation Change

As indicated in Table 1, the port block allocation change event is indicated by MSG-ID set to "PtAlloc". The associated SD-ELEMENT is tagged by SD-ID "NATPBlk". The contents of the NATPBlk SD-ELEMENT are shown in Table 5. The requirements for these contents are derived from the description in Section 3.3.

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

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

As in the example in the previous section example, consider a DS-Lite AFTR [RFC6333] incorporating a PCP server, where PCP is used to

obtain an external address binding and a port range. See [I-D.pcp-port-set] for the port set part of this operation.

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

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

Character count: about 155.

5.3.4. Address Exhaustion Event

As indicated in Table 1, the address exhaustion event is indicated by MSG-ID set to "AddrEx". The associated SD-ELEMENT is tagged by SD-ID "NATAddrEx". The contents of the NATAddrEx SD-ELEMENT are shown in Table 6. The requirements for these contents are derived from the description in Section 3.4.

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

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

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

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

Character count: about 120.

5.3.5. NAT Port Exhaustion

As indicated in Table 1, the port exhaustion event is indicated by MSG-ID set to "PortEx". The associated SD-ELEMENT is tagged by SD-ID "NATPEx". The contents of the NATPEx SD-ELEMENT are shown in Table 7. The requirements for these contents are derived from the description in Section 3.5.

PARAM-NAME	Description	Requirement
DevTyp	Section 5.2.3	OPTIONAL
DevID	Section 5.2.2	OPTIONAL
PostS4	Section 5.2.4	MANDATORY
Proto	Section 5.2.7	MANDATORY

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

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

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

Character count: about 110.

5.3.6. Quota Exceeded

As indicated in Table 1, the quota exceeded event is indicated by MSG-ID set to "Quota". The associated SD-ELEMENT is tagged by SD-ID "NATQEx". The contents of the NATQEx SD-ELEMENT are shown in Table 8. The requirements for these contents are derived from the description in Section 3.6.

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

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

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

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

SiteID="A2E0:62"]

Character count: about 135.

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

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

Character count: about 115.

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

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

Character count: about 100.

5.3.7. Invalid Port Detected

As indicated in Table 1, the invalid port detected event is indicated by MSG-ID set to "InvPort". The associated SD-ELEMENT is tagged by SD-ID "NATInvP". The contents of the NATInvP SD-ELEMENT are shown in Table 9. The requirements for these contents are derived from the description in Section 3.7.

PARAM-NAME	Description	Requirement
DevID	Section 5.2.2	OPTIONAL
SiteID	Section 5.2.11	MANDATORY
PSID	Section 5.2.9	MANDATORY

Table 9: Contents Of the SD-ELEMENT Section For Logging the Invalid Port detected Event

Example: Lightweight 4over6 BR configured with PSID 15 for the given subscriber site.

<83>1 2013-05-07T15:27:49.603Z yourd137mzmhow.example.net
NAT 68 InvPort [NATInvP SiteID="5A27:876E" PSID="15"]

Character count: about 110.

6. IANA Considerations

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

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

	VLANid	OPTIONAL	RFCxxxx
	VRFid	OPTIONAL	RFCxxxx
----	----	----	----
NATInvP		OPTIONAL	RFCxxxx
	DevID	OPTIONAL	RFCxxxx
	SiteID	MANDATORY	RFCxxxx
	PSID	MANDATORY	RFCxxxx

Table 10: NAT-Related STRUCTERED-DATA Registrations

IANA is further requested to establish a new registry entitled "syslog NAT Types" within the "syslog Parameters" registry. The initial values for this registry are shown in Table 11. New values may be added following the criterion of IETF Review.

Value	Description	Reference
44	NAT44	RFCxxxx
64	NAT64	RFCxxxx
AFTR	DS-Lite AFTR [RFC6333]	RFCxxxx
BR	Lightweight 4over6 or MAP-E BR	RFCxxxx

Table 11: syslog NAT Type Values

7. Security Considerations

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

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

8. References

8.1. 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.

8.2. Informative References

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

[I-D.pcp-port-set] Sun, Q., Boucadair, M., Sivakumar, S., Zhou, C., Tsou, T., and S. Perreault, "Port Control Protocol (PCP) Extension for Port Set Allocation (Work in progress)", March 2013.

[I-D.softwire-lw4over6] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture (Work in progress)", April 2013.

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

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

[I-D.softwire-unified-cpe]
Boucadair, M. and I. Farrer, "Unified IPv4-in-IPv6 Softwire CPE (Work in progress)", March 2013.

[I-D.tsou-behave-natx4-log-reduction]
Tsou, T., Li, W., and T. Taylor, "Port Management To Reduce Logging In Large-Scale NATs (Work in progress)", May 2013.

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

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

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

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

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

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

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

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

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

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

Authors' Addresses

Zhonghua Chen
China Telecom
P.R. China

Email: 18918588897@189.cn

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

Email: cathy.zhou@huawei.com

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

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

T. Taylor
Huawei Technologies
Ottawa
Canada

Email: tom.taylor.stds@gmail.com