



External Logging

External logging configures the export and logging of the NAT table entries, private bindings that are associated with a particular global IP port address, and to use Netflow to export the NAT table entries.

- [Bulk Port Allocation, on page 1](#)
- [Session logging, on page 2](#)
- [\(Only NAT44\) Syslog Logging, on page 2](#)
- [Reliable Log Transfer, on page 19](#)
- [Frequently Asked Questions \(FAQs\), on page 21](#)

Bulk Port Allocation

The creation and deletion of NAT sessions lead to creation of logs. If logs of all such translations are stored, then a huge volume of data is created. This data is stored on a NetFlow or a Syslog collector. To reduce the volume of this data, a block of ports is allocated. If bulk port allocation is enabled, as soon as a subscriber creates the first session, a number of contiguous external ports are allocated. To indicate this allocation, a bulk allocation message is created in the log.



Note The bulk allocation message is created only during the first session. Rest of the sessions use one of the allocated ports. Hence no logs are created for them.

A bulk delete message is created in the log when the subscriber deletes all the sessions that are using the allocated ports.

Another pool of ports is allocated only if the number of simultaneous sessions is more than N where N is the size of the bulk allocation. The size of the pool can be configured from the CLI.

Restrictions for Bulk Port Allocation

The restrictions for bulk port allocation are as follows:

- The value for the size of bulk allocation can be 16, 32, 64, 128, 256, 512, 1024, 2048 and 4096. For optimum results, it is recommended that you set this size to half of the port limit.
- If the size of bulk allocation is changed, then all the current dynamic transactions will be deleted. Hence it is advisable to change the bulk port allocation size (only if necessary) during a maintenance window.

- The port numbers below the value of dynamic-port-range start value (which is 1024 by default), are not allocated in bulk.
- The algorithm that is used to allocate a public address to a user remains the same.
- When bulk allocation is enabled, session logging is not available.
- When bulk allocation is enabled, the translation record will not contain information about L4 protocol.
- Bulk port allocation features is not supported in NAT64 stateful application.

Session logging

In general, NAT translation entries contain information about private source IP, port and translated public IP and port. However, there could be cases when the destination IP address (public IP address) and port may also be needed. In such cases, session logging has to be enabled so that Netflow or Syslog translation records include these values as well.

(Only NAT44) Syslog Logging

Perform the following tasks to configure Syslog Logging for NAT table entries.

Restrictions for Syslog

The restrictions for syslog are as follows:

- Syslog is supported over UDP only.
- Syslog is supported in ASCII format only.
- You cannot log onto multiple collectors or relay agents.
- All the messages comply to RFC 3954 except for the timestamp format. Timestamp is represented in a simpler way as explained later in this section.
- Syslog shall be supported for DS-Lite and NAT444 as of now. Support for NAT64 is not yet available.

Syslog Message Format

In general, the syslog message is made up of header, structured data, and msg fields. However, in the CGv6 applications, the structured data is not used.

Header

The header fields shall be as per the RFC 5424. Fields shall be separated by ' ' (blank space) as per the RFC.

The header consists of the following fields:

Field	Description
Priority	<ul style="list-style-type: none"> The priority value represents both the facility and severity. Ensure that the severity code is set to Informational for all the messages at value 6.
Version	<ul style="list-style-type: none"> This field denotes the version of the specification of the syslog protocol. In CGv6 application, the version value is set to 1.
Timestamp	<ul style="list-style-type: none"> This field is needed to trace the port usage. The format is <year> <mon> <day> <hh:mm:ss>. Ensure that the syslog collector converts the time to local time whenever needed. <p>Note The timestamp is always reported in GMT/UTC irrespective of the time zone configured on the device.</p>
Hostname	<ul style="list-style-type: none"> This field is used to identify the device that sent the syslog message. While configuring the syslog server, ensure that the host name does not exceed 31 characters. The default value for the host name is '-'.
App name and PROC ID	These fields are not included. In ASCII format, '-' is included for these fields.
MSG ID	<ul style="list-style-type: none"> This field identifies the type of the syslog message. In the ASCII format, the values for NAT44 and DS Lite messages are NAT44 and DS LITE respectively.

Structured Data

It is not used.

MSG

This field consists of the information about the NAT44 or DS Lite events. In a single UDP packet, there could be one or more MSG fields each enclosed in [] brackets. The MSG field has many sub fields as it has a common structure across different records (for both NAT44 and DS Lite). Note, that, depending on the event, some of

the fields may not be applicable. For example, fields such as 'Original Source IPv6' address are not applicable for all NAT44 events. In such cases, the inapplicable fields will be replaced by '-'.

The syntax of the MSG part is as follows:

[EventName <L4> <Original Source IP> <Inside VRF Name> <Original Source IPv6> <Translated Source IP> <Original Port> <Translated First Source Port> <Translated Last Source Port> <Destination IP> <Destination Port>]

The descriptions of the fields in this format are as follows:

Field	Description
EventName	<p>Select any one of the values for EventName from the following based on the event:</p> <ul style="list-style-type: none"> • UserbasedA: User-based port assignment • SessionbasedA: Session-based port assignment • SessionbasedAD: Session-based port assignment with destination information <p>Note: SessionbasedAD is used only if session logging is enabled. Also, session-logging and bulk port allocation are mutually exclusive.</p> <ul style="list-style-type: none"> • UserbasedW: User-based port withdrawal • SessionbasedW: Session-based port withdrawal • SessionbasedWD: Session-based port withdrawal with destination information • Portblockrunout: Ports exhausted
L4	<p>Specifies the identifier for the transport layer protocol. Select any one of the values for L4 from the following:</p> <ul style="list-style-type: none"> • 1 for ICMP • 6 for TCP • 17 for UDP • 47 for GRE
Original Source IP	Specify the private IPv4 address.
Inside VRF Name	<p>The Inside vrf is essential to identify the subscriber. Even though multiple subscribers connected to the router might have the same source IP, they might be placed in different VRFs. Hence the VRF name and the original source IP together helps to identify a subscriber. Ensure that the VRF name is not more than 32 characters in length.</p>

Field	Description
Original Source IPv6	Specifies the IPv6 source address of the tunnel in case of DS Lite.
Translated Source IP	Specifies the public IPv4 address post translation.
Original Port	Specifies the source port number before translation. This is not applicable for the UserbasedA and UserbasedW events.
Translated First Source Port	Specifies the first source port after translation.
Translated Last Source Port	Specifies the last source port after translation. This is applicable only for the UserbasedA and UserbasedW events.
Destination IP	Specifies the destination IP recorded in the syslogs for the SessionbasedAD and SessionbasedWD events.
Destination Port	Specifies the destination port recorded in the syslogs for the SessionbasedAD and SessionbasedWD events.

Let us look at an example for NAT444 user-based UDP port translation mapping:

```
[UserbasedA - 10.0.0.1 Broadband - 100.1.1.1 - 2048 3071 - -]
```

The description for this example is as follows:

Value	Description
UserbasedA	Event Name
10.0.0.1	Original Source IP
Broadband	Inside VRF name
100.1.1.1	Translated Source IP
2048	Translated First Source Port
3071	Translated Last Source Port



- Note** The number of MSG fields in an UDP packet are determined by the following factors:
- The space available in the UDP packet depends on MTU.
 - The translation events pertaining to MSG records in a given packet must have happened within a second (starting from the time at which the first event of that packet happened).

Netflow v9 Support

The NAT64 stateful, NAT44, and DS Lite features support Netflow for logging of the translation records.. The Netflow uses binary format and hence requires software to parse and present the translation records. However, for the same reason, Netflow requires lesser space than Syslog to preserve the logs.

Considerations

The considerations for NetFlow are as follows:

- NetFlow V9 is supported over UDP.
- You cannot log onto multiple collectors or relay agents.
- All the messages comply to RFC 3954.

NetFlow Record Format

As NetFlow V9 is based on templates, the record format contains a packet header and templates or data records based on templates.

Header

All the fields of the header follow the format prescribed in RFC 3954. The source ID field is composed of the IPv4 address of ServiceInfra interface (of the card) and specific CPU-core that is generating the record. The collector device can use the combination of the sourceIPv4Address field plus the Source ID field to associate an incoming NetFlow export packet with a unique instance of NetFlow on a particular device.

Templates for NAT44

The templates are defined and used for logging various NAT64 stateful, NAT44 and DS Lite events as follows. The templates may change in future software releases. Hence it is advised that the Netflow collector software is designed to understand the templates as distributed by the router and accordingly parse the records.

Options Templates

The translation entries consist of VRF IDs which might be incomprehensible to a user. To simplify this process, the CGv6 applications send the options templates along with the data templates.

Options template is a special type of data record that indicates the format of option data related to the process of NetFlow. The options data consist of the mapping between VRF IDs and VRF names. By parsing and using this data, the NetFlow collectors can modify the translation entries by adding VRF names instead of VRF IDs.

The value for the Template ID of options template is 1 where as the value of the Template ID for data template is 0. For more information on Options template, see RFC3954.

Events

The events and the corresponding template details are described in the following table:

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
Nat444 translation create event	256	Disabled	Disabled	ingressVRFID	234	4	ID of the Ingress VRF
				egressVRFID	235	4	ID of the Egress VRF
				sourceIPv4Address (pre-NAT)	8	4	Original Source IPv4 address
				postNATSourceIPv4 Address	225	4	Post NAT (outside) source IPV4 address
				sourceTransportPort (pre NAT)	7	2	Original source port
				sourceTransportPort	7	2	Post NAT (inside) source port
				protocolIdentifier	4	1	L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
Nat444 session create event - session based (with destination)	271	Disabled	Enabled	ingressVRFID	234	4	ID of the Ingress VRF
				egressVRFID	235	4	ID of the Egress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				sourceTransportPort	7	2	Original Source Port
				sourceTransportPort	7	2	Post NAT (inside) source port
				destinationIPv4Address	12	4	Destination IP address
				destinationTransportPort	11	2	Destination port
				protocolIdentifier	4	1	L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
Nat444 translation create event - user based	265	Enabled	Disabled	ingressVRFID	234	4	ID of the Ingress VRF
				egressVRFID	235	4	ID of the Egress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				postNATPortBlockStart	361	2	Start of Post NAT (inside) source port block
				postNATPortBlockEnd	362	2	End of Post NAT source port block

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
Nat444 translation delete event	257	Disabled		ingressVRFID	234	4	ID of the Ingress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				sourceTransportPort	7	2	Original source port
				protocolIdentifier	4	1	L4 protocol identifier
Nat444 session delete event - session based (with destination)	272	Disabled	Enabled	ingressVRFID	234	4	ID of the Ingress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				destinationIPv4Address	12	4	Destination IP address
				sourceTransportPort	7	2	Post NAT (translated) source port
				destinationTransportPort	11	2	Destination port
				protocolIdentifier	4	1	L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
Nat444 translation delete event - user based	266	Disabled	Disabled	ingressVRFID	234	4	ID of the Ingress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				postNATPortBlockStart	361	2	Start of Post NAT (and) source port block. Note this is not defined by IANA yet.
DS-Lite translation create event	267	Disabled	Disabled	ingressVRFID	234	4	ID of the Ingress VRF
				egressVRFID	235	4	ID of the Egress VRF

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
				Pre NAT Source IPv4 Address	8	4	Original source IPV4 address. This field is valid only when source is enabled. Else, it will be reported as 0
				Pre NAT Source IPv6 Address	27	16	IPv6 address of the B4 element (Tunnel source)
				postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				sourceTransportPort	7	2	Original source port
				sourceTransportPort	227	2	Post NAT (inside) source port

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
DS-Lite session create event - session based (with destination)	273	Disabled	Enabled	ingressVRFID	234	4	ID of the Ingress VRF
				egressVRFID	235	4	ID of the Egress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				sourceIPv6Address	27	16	IPv6 address of the B4 element (Tunnel source)
		Enabled	Disabled	postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				sourceTransportPort	7	2	Original source port
				sourceTransportPort	227	2	Post NAT (inside) source port
				destinationIPv4Address	12	4	Destination IP address
				destinationTransportPort	11	2	Destination port
				protocolIdentifier	4	1	L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
DS-Lite translation create event - user based	269	Enabled	Disabled	ingressVRFID	234	4	ID of the Ingress VRF
				egressVRFID	235	4	ID of the Egress VRF
				sourceIPv4Address	8	4	Original source IPV4 address. This field is valid only when source is enabled. Else, it will be reported as 0
				sourceIPv6Address	27	16	IPv6 address of the B4 element (Tunnel source)

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
DS-Lite translation create event - user based				postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				postNATPortBlockStart	361	2	Start of Post NAT (inside) source port block
				postNATPortBlockEnd	362	2	End of Post NAT source port block
DS-Lite translation delete event	270	Disabled	Disabled	ingressVRFID	234	4	ID of the Ingress VRF
				sourceIPv4Address			Original source IPV4 address
				sourceIPv6Address			IPv6 address of the B4 element (Tunnel source)
				sourceTransportPort			Original source port
				protocolIdentifier			L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
DS-Lite session delete event - session based (with destination)				ingressVRFID	234	4	ID of the Ingress VRF
				sourceIPv4Address	8	4	Original source IPV4 address
				sourceIPv6Address	27	16	IPv6 address of the B4 element (Tunnel source)
				sourceTransportPort	7	2	Original source port
				protocolIdentifier	4	1	L4 protocol identifier
DS-Lite translation delete event - user based	270	Disabled	Disabled	ingressVRFID	234	4	ingressVRFID
				sourceIPv4Address	8	4	Original source IPV4 address
				sourceIPv6Address	27	16	IPv6 address of the B4 element (Tunnel source)
				postNATPortBlockStart	361	2	Start of Post NAT (translated) source port block

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
NAT64 stateful translation create event	258	Disabled	Disabled	sourceIPv6Address	27	16	Source IPv6 address
				postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				sourceTransportPort	7	2	Original source port
				sourceTransportPort	227	2	Post NAT (inside) source port
				protocolIdentifier	4	1	L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
NAT64 stateful session create event - session based (with destination)	260	Disabled	Enabled	sourceIPv6Address	27	16	Source IPv6 address (pre translation)
				postNATSourceIPv4Address	225	4	Post NAT (outside) source IPV4 address
				destinationIPv6Address	28	16	Destination IPv6 address (pre translation)
				Post translation Destination IP address	226	4	Destination IPv4 address (post translation)
				sourceTransportPort	7	2	Original source port
				sourceTransportPort	7	2	Post NAT (translated) source port
				destinationTransportPort	11	2	Destination port
				protocolIdentifier	4	1	L4 protocol identifier

Event	Template ID	Bulk Port Allocation	Destination/Session Logging	Field Name	IANA IPFIX ID	Size in bytes	Description
NAT64 translation delete event	259	Disabled	Disabled	sourceIPv6Address	27	16	IPv6 address of the B4 element (Tunnel source)
				sourceTransportPort	7	2	Original source port
				protocolIdentifier	4	1	L4 protocol identifier
NAT64 stateful session delete event - session based (with destination)	261	Disabled	Enabled	sourceIPv6Address	27	16	IPv6 address of the B4 element (Tunnel source)
				destinationIPv6Address	28	16	Destination IPv6 address (pre translation)
				sourceTransportPort	7	2	Original source port
				destinationTransportPort	11	2	Destination port
				protocolIdentifier	4	1	L4 protocol identifier

Reliable Log Transfer

The VSM Line Cards support CGN application with NAT 44 feature. NAT 44 feature performs IPv4 NAT operations. The CGN applications based on server configuration generates NetFlow and Syslog records that contain critical NAT information. If an external NetFlow server is configured for storage and retrieval of NAT information, the information is transferred to the server using UDP connection.

The CGN application VM running on VSM also supports transfer of records to external NetFlow servers using TCP. This provides reliable log transfer. Use the *protocol* command to define the protocol to be used for the transfer.

Limitations

- Netflow and syslog records may be lost during any reload of VSM or restart of IOS XR VM or process.
- A maximum of two NetFlow or Syslog servers are supported.
- Server IP address along with the port number defines one server.
- Same TCP server can be configured on any number of VRFs and is considered as one TCP server.
- Loss of TCP connection may result in loss of NetFlow records.
- Bulk port allocation needs to be enabled with a minimum allocation value of 256 to prevent loss of NAT 44 records, when Syslog and Netflow records are sent using TCP protocol.

Reliable Log Transfer Configuration: Examples

Configure NetFlow records' transport using TCP protocol when the TCP server is in the default global VRF instance:

```
service cg n1
service-location preferred-active 0/3/CPU0
service-type nat44 nat1
inside-vrf insidevrf1
map outside-vrf outsidevrf1 outsideServiceApp ServiceApp2 address-pool 100.1.1.0/24
external-logging netflow version 9
server
address 45.1.1.1 port 65535
protocol tcp
```

Configure NetFlow records' transport using TCP protocol when the TCP server is in a non-default VRF instance:

```
service cg n1
service-location preferred-active 0/3/CPU0
service-type nat44 nat1
inside-vrf insidevrf1
map outside-vrf outsidevrf1 outsideServiceApp ServiceApp2 address-pool 100.1.1.0/24
external-logging netflow version 9
server
address 55.1.1.1 port 4500
protocol tcp vrf netflow-srv-vrf
```

Disabling NetFlow records' transport:

```
service cg n1
service-type nat44 nat1
inside-vrf insidevrf1
no external-logging netflow version 9
```

Configure Syslog records' transport using TCP protocol when the TCP server is in the default global VRF instance:

```

service cgn cgn1
service-location preferred-active 0/3/CPU0
service-type nat44 nat1
inside-vrf insidevrf1
map outside-vrf outsidevrf1 outsideServiceApp ServiceApp2 address-pool 100.1.1.0/24
external-logging syslog
server
address 45.1.1.1 port 3500
protocol tcp

```

Configure Syslog records' transport using TCP protocol when the TCP server is in an non-default VRF instance:

```

service cgn cgn1
service-location preferred-active 0/3/CPU0
service-type nat44 nat1
inside-vrf insidevrf1
map outside-vrf outsidevrf1 outsideServiceApp ServiceApp2 address-pool 100.1.1.0/24
external-logging syslog
server
address 55.1.1.1 port 4500
protocol tcp vrf syslog-srv-vrf

```

Disabling Syslog records' transport:

```

service cgn cgn1
service-type nat44 nat1
inside-vrf insidevrf1
no external-logging syslog

```

Frequently Asked Questions (FAQs)

This section provides answers to the following frequently asked questions on external logging.

Q: How to trace a subscriber by using the NAT logs?

A: In order to trace a subscriber, you should know the public source IP address (post NAT source address), post NAT source port, protocol, and the time of usage. With these parameters, the steps to trace a subscriber are as follows:

1. Search for the create event that has the matching public IP address, post NAT Source IP address (postNATSourceIPv4Address) and protocol, egress VRF ID/Name and the time of the usage. Ensure that the time of the create-event is the same or earlier than the time of usage reported. You may not find the protocol entry or the exact post NAT source port in the logs if bulk allocation is enabled. In such cases, find the create-event whose **Post NAT Port Block Start** and **Post NAT Port Block End** values include the post NAT source port. The **Pre NAT source IP address** along with the corresponding **ingress VRF ID/Name** will identify the subscriber.
2. The corresponding delete record may be found optionally to confirm that the subscriber was using the specified public IP and port during the time of the reported usage.

Q: The Netflow records provide VRF IDs for ingress and egress VRFs. How will I know the VRF names?

A: The following are the two ways to find the VRF name from the VRF ID.

1. Use the command `show rsi vrf-id <vrf-id>` on the Router console to find VRF-ID to VRF-NAME associations.

2. The CGv6 applications periodically send out option templates containing the VRF-ID to VRF-NAME mapping. The Netflow collector software presents the information with VRF-Names rather than VRF IDs.

Q: Does the time format in Syslog or Netflow account for Day light saving?

A: The Syslog and Netflow formats report time corresponding to GMT/UTC. The Netflow header contains the time in seconds that elapsed since EPOCH whereas the Syslog header contains time in human readable formats. In both cases, the day light saving is not accounted. The Netflow/Syslog collectors have to make that adjustments if needed.

Q: Since the Netflow and Syslog use UDP, how can we know if a packet containing translation record was lost?

A: The Netflow header contains a field called Sequence Number. This number indicates the count of the packet coming from each Source ID. The Netflow collector traces the Sequence Number pertaining to each unique Source ID. The sequence numbers should be increased by one for each packet sent out by the Source. If the collector ever receives two successive packets with the same Source ID, but with a Sequence number difference of more than 1, it indicates a packet loss. However, currently, no such mechanism exists for Syslog.

Q: What is the use of session-logging?

A: Session logging includes destination IP and port number as well. Though this information is not directly useful in tracing the subscriber, in some cases, this information may be useful or may be mandated by the legal authorities. There are cases where, legal authorities may not have the post NAT source 'port', however may know the destination IP address (and optionally destination port, such as IP address and port of an e-mail server). In the absence of post NAT source port information, a list of subscribers who used the specified public IP during that time may have to be pruned further based on the destination IP and port information.

Q: How does the bulk port allocation reduce data volume of translation logs?

A: With bulk port allocation, subscribers are allocated a range of contiguous ports on a public IP. Quite often, a subscriber will need more ports than just one. Especially AJAX based web pages and other web applications simultaneously open several ports. In such cases, pre-allocated ports are used and only one log entry is made that specifies the range of ports allocated to the user. Hence, bulk port allocation significantly reduces log data volume and hence the demand on storage space needed for the translation logs.

Q: What else can be done to reduce log data volume?

A: Predefined NAT is an option that can be used to eliminate the logging altogether. The Predefined NAT translates private IP address to public IP address and a certain port range by using an algorithm. Hence there is no need to keep track of NAT entries.