



Network Address Translation Overview

This chapter provides an overview of Network Address Translation (NAT) in-line service feature.

The following topics are covered in this chapter:

- [NAT Overview, on page 1](#)
- [How NAT Works, on page 30](#)

NAT Overview

This section provides an overview of the NAT in-line service feature.

NAT translates non-routable private IP address(es) to routable public IP address(es) from a pool of public IP addresses that have been designated for NAT. This enables to conserve on the number of public IP addresses required to communicate with external networks, and ensures security as the IP address scheme for the internal network is masked from external hosts, and each outgoing and incoming packet goes through the translation process.

The NAT in-line service works in conjunction with the following products:

- GGSN
- HA
- PDSN
- P-GW
- SaMOG

NAT works by inspecting both incoming and outgoing IP datagrams and, as needed, modifying the source IP address and port number in the IP header to reflect the configured NAT address mapping for outgoing datagrams. The reverse NAT translation is applied to incoming datagrams.

NAT can be used to perform address translation for simple IP and mobile IP. NAT can be selectively applied/denied to different flows (5-tuple connections) originating from subscribers based on the flows' L3/L4 characteristics—Source-IP, Source-Port, Destination-IP, Destination-Port, and Protocol.



Important

NAT works only on flows originating internally. Bi-directional NAT is not supported.

**Important**

NAT is supported only for TCP, UDP, and ICMP flows. For other flows NAT is bypassed. For GRE flows, NAT is supported only if the PPTP ALG is configured. For more information on ALGs, please refer to the *NAT Application Level Gateway* section.

**Important**

In 14.1 and earlier releases: If a subscriber is assigned with a public IP address, NAT is not applied. For 15.0 and later releases, NAT can be applied for private and public addresses if the IP pool is configured with the **skip-nat-subscriber-ip-check** CLI option.

**Important**

To get NATed, the private IP addresses assigned to subscribers must be from the following ranges: Class A 10.0.0.0 – 10.255.255.255, Class B 172.16.0.0 – 172.31.255.255, and Class C 192.168.0.0 – 192.168.255.255, and 100.64.0.0/10 as per RFC 6598.

As per a new implementation, NAT can now be enabled or disabled irrespective of whether the IP assigned is a private or public IP by enabling a CLI option in IP pool. On enabling this option, the private IP check for the corresponding pool will be skipped and NAT will be enabled (if configured) for this pool although it is a public pool. Refer to the *Configuring One-to-One NAT IP Pools/NAT IP Pool Groups* section in the *NAT Configuration* chapter for more information.

NAT supports the following mappings:

Once a flow is marked to use a specific NAT IP address the same NAT IP address is used for all packets originating on that flow. The NAT IP address is released only when all flows and subscribers associated with it are released.

When all NAT IP addresses are in use, and a subscriber with a private IP address fails to get a NAT IP address for a specific flow, that specific flow will not be allowed and will fail.

All downlink—inbound from external networks—IP packets that do not match one of the existing NAT bindings are discarded by the system.

Qualified Platforms

NAT is a StarOS in-line service application that runs on Cisco ASR 5500 and virtualized platforms. For additional platform information, refer to the appropriate *System Administration Guide* and/or contact your Cisco account representative.

License Requirements

The NAT is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

NAT Realms

A NAT realm is a pool of unique public IP addresses available for translation from private source IP addresses. IP addresses in a NAT IP pool are contiguous, and assignable as a subnet or a range that constitutes less than

an entire subnet. IP addresses configured in NAT IP pools within a context must not overlap. At any time, within a context, a NAT IP address must be configured in any one NAT IP pool. IP addresses can be added to a NAT IP pool as a range of IP addresses. Based on the chosen port chunk-size, the number of subscribers that can be shared per IP varies.

The minimum number of public IP addresses that must be allocated to each NAT IP pool must be equal to the number of Session Managers (SessMgrs) available on the system. Theoretically, the number of SessMgrs that can be brought up is 384. However, the number of SessMgrs can vary based on the cards on the system.

Up to 2000 unique “IP pools + NAT IP pools” can be configured per context. A maximum of twenty NAT IP pools/NAT IP pool groups can be configured in a Firewall-and-NAT policy. At any time a subscriber can be associated with a maximum of three different NAT IP pools/NAT IP pool groups and can have NATED flows on three different NAT IP addresses at the same time.

In 20 and later releases, each many-to-one NAT realm can support multiple NAT IP addresses for the same NAT realm for a given subscriber. If no ports are available for a given NAT IP, then instead of dropping packets, another NAT IP will be requested for the same NAT realm as long as the maximum number of port chunks configured is not reached. The number of NAT IPs that can be allocated for a given NAT realm for a particular subscriber is limited to a maximum of three IPs. This is applicable only to many-to-one NAT realms. Refer to the *Configuring IP address allocation for NAT realm* section in the *NAT Configuration* chapter for more information on enabling and disabling this feature.

Allocation of NAT IP addresses in NAT IP pools to subscriber traffic is based on the L3/L4 characteristics—IP addresses, ports, and protocol—of the subscriber flows. It is possible to configure the system to perform or not perform NAT based on one or more L3/L4 parameters. This feature is also known as Target-based NAT. For more information, see the *Target-based NAT Configuration* section.

The following table lists the minimum requirements for configuring NAT pools.

Card Type	No. of SMGRs	No. of NAT Translations per SMGR	No. of Active Cards	Total no. of NAT Translations
PSC	7	317000	12	26628000 (26.628M)
PSC2	16	270000	12	51840000 (51.84M)
PSC3	24	400000	12	115200000 (115.2M)

NAT IP pools have the following configurable parameters. These parameters are applicable to all IP addresses in a NAT IP pool.

- NAT IP Address Allocation Mode: Specifies when to allocate a NAT IP address to a subscriber; either at call setup or during data flow based on the allocation mode.
 - Not-on-demand Allocation Mode: This is the default mode. In this mode, the NAT IP address is allocated to the subscriber at call setup. If there are three NAT IP pools/NAT IP pool groups configured in the subscriber's Firewall-and-NAT policy, the subscriber is allocated three NAT IP addresses, one from each NAT IP pool/NAT IP pool group during call setup. If NAT IP address is not available for any of the pools, then the subscriber call is dropped.
 - On-demand Allocation Mode: In this mode NAT resources are assigned and allocated dynamically based on subscriber flows. The NAT IP address is allocated to the subscriber when the data traffic flows in and not at call setup.

In case of on-demand pools, since the NAT IP address is not allocated to the subscriber at call setup, the subscriber may not have a NAT IP address allocated when the first packet is received. Until the

successful allocation of a NAT IP address, based on the configuration, the packets can either be buffered or dropped. Once a free NAT IP address is available, it is allocated to the subscriber to be used for flows matching the pool.

For On-demand NAT realms, the subscribers can be filtered based on NAT IP usage time to find out how long (in seconds) the subscriber has been using the assigned NAT IP.

- **NAT Binding Timer:** Specifies the timeout period, in seconds, to deallocate NAT resources that were allocated to subscriber flows. When a subscriber flow stops the timer starts counting down, and on expiry the NAT resources are deallocated to be made available for other subscriber flows.
 - In one-to-one allocation, for a given NAT IP address, the NAT Binding Timer starts counting down when there are no active flows using that NAT IP address. When the NAT Binding Timer expires, the NAT IP address gets deallocated.
 - In many-to-one allocation, wherein subscribers are allocated port-chunks rather than individual ports, as long as a port-chunk is allocated to a subscriber, all ports from that port-chunk are reserved for that subscriber. When all flows using ports from that port-chunk get timed out/cleared, the NAT Binding Timer starts counting down. If any new flows come up before the NAT Binding Timer expires, ports are once again allocated from that port-chunk, and the NAT Binding Timer gets cancelled. As long as there are active flows using the port-chunk it cannot be deallocated. But, if no new flows come and the NAT Binding Timer expires, the port-chunk gets deallocated. In the case of on-demand NAT, if it is the last port-chunk for the NAT IP address, on NAT Binding Timer expiry, the NAT IP address gets deallocated along with the last port-chunk.
- **Maximum Users per NAT IP Address:** Applicable only to many-to-one NAT IP pools. Specifies the maximum number of subscribers sharing one NAT IP address.

In 18 and earlier releases, a maximum number of 2016 subscribers can be configured per NAT IP address.

In 19 and later releases, a maximum number of 8064 subscribers can be configured per NAT IP address.

- **Port Chunk Size:** Applicable only to many-to-one NAT IP pools. Specifies the block size of contiguous ports to be assigned to a many-to-one NAT subscriber.

In 18 and earlier releases, the minimum supported port chunk size was 32 and the chunk size was configurable in multiples of 32, that is, 32, 64, 96, and so on. This number has to be divisible by 32 up to a maximum of 32,256.

In 19 and later releases, the minimum port chunk size configurable is reduced to 8 and the chunk size can be configured in multiples of 8 starting 8, 16, 24, 32, and so on. The default port chunk size is 32. If no chunk size is configured, then the chunk size is calculated by dividing the entire NAT port range by the number of users per NAT IP and then rounding off to the nearest multiple of 32.

- **Maximum Port-chunks per User:** Applicable only to many-to-one NAT IP pools. Specifies the maximum number of port-chunks allowed for an individual subscriber from the same NAT IP address. This will limit subscribers from dominating all the available ports in a many-to-one NAT IP.

In 18 and earlier releases, a maximum number of 2016 port chunks can be configured per subscriber.

In 19 and later releases, the maximum number of port chunks that can be configured per subscriber is 8064.

Consider a case where a single TCP flow is active in a port-chunk. When this connection gets cleared, the TCP NAT port goes to Time Wait state. Since it is the last flow of the port-chunk, the NAT Binding Timer also gets started. Assume NAT Binding Timer \geq TCP 2MSL Timer. Once the 2MSL Timer expires, the TCP port would go to Free state. However, the NAT Binding Timer keeps running. On NAT

Binding Timer expiry, the port-chunk is deallocated. If this was the last port-chunk for that subscriber, the NAT IP address is also deallocated along with this port-chunk.

In case NAT Binding Timer < TCP 2MSL Timer, at NAT Binding Timer expiry, the TCP port is forcefully moved to Free state from Time Wait state and the port-chunk deallocated.

- Port Chunk Thresholds: Applicable only to many-to-one NAT IP pools. Specifies threshold in terms of percentage of allocated port-chunks against total port-chunks available. Once the threshold is reached, new subscribers will not be allocated the same NAT IP address.
- Packet Drop Thresholds: Specifies threshold in terms of percentage of NAT packet drops.
- AAA Binding Update Message Required: Applicable only to one-to-one NAT IP pools. Enables AAA binding messages for one-to-one NAT IP pools. This is not supported for many-to-one NAT IP pools.
- Alert Thresholds: Threshold limits can be specified to trigger alarms for NAT IP pools for pool-used, pool-free, pool-hold, and pool-release cases.
- SRP-Activate: Applicable to both one-to-one and many-to-one NAT IP pools. When configured, the NAT IP pool will become usable only when the SRP state is active.

Network broadcast is supported for NAT pools and ordinary pools. The Busyout feature is also supported for NAT pools and ordinary pools.

NAT IP Pool Groups

Similar NAT IP pools can be grouped into NAT IP pool groups. This enables to bind discontinuous IP address blocks in individual NAT IP pools to a single NAT IP pool group.

When configuring a NAT IP pool group, note that only those NAT IP pools that have similar characteristics can be grouped together. The similarity is determined by the NAT IP pool Type (One-to-One / Many-to-One), users configured per NAT IP address (applicable only to many-to-one NAT IP pools), NAT IP Address Allocation Mode (On Not-on-demand), and Port Chunk Size (applicable only to many-to-one NAT IP pools) parameters. Dissimilar NAT IP pools cannot be grouped together.

It is recommended that all the NAT IP pools in a NAT IP pool group be configured with the same values for the other parameters, so that the NAT behavior is predictable across all NAT IP pools in that NAT IP pool group.

The NAT IP pool from which a NAT IP address is assigned will determine the actual values to use for all parameters.

It is recommended that in a Firewall-and-NAT policy all the realms configured either be NAT IP pools or NAT IP pool groups. If both NAT IP pool(s) and NAT IP pool group(s) are configured, ensure that none of the NAT IP pool(s) are also included in the NAT IP pool group.

NAT IP Address Allocation and Deallocation

Cisco System's implementation of NAT is Endpoint-independent Mapping, wherein NAT reuses the same NAT source port mapping for subsequent packets sent from the same private IP address and port, and with the same protocol to any public destination host IP address and port.

That is, all flows coming from the subscriber for the current session with the same protocol and same source IP address and source port (X:x) would get the same NAT IP address and NAT port (X:x) irrespective of the destination IP address and port. NAT will not allow any inbound packets to the NAT IP address and NAT

port (X:x) from an external host IP address and host port (Y:y), unless the internal host (MS) had previously sent a packet of the same protocol type to that external IP address and Port (Y:y). However, this behavior changes if NAT ALG is enabled. The ALG creates pin holes / dynamic routes in the NAT and allows downlink packets that match the pin holes / dynamic routes towards the internal host (MS) given that there was already a parent connection from MS towards the external host.

The advantage of endpoint-independent mapping is that applications are unaffected by NAT translations.

Inbound connection to the NAT IP address can be allowed in one-to-one pools based on configuration.

NAT IP Address Allocation

The NAT IP address is allocated based on the following parameters:

- **Maximum Users per NAT IP Address:** The maximum number of subscribers sharing a NAT IP address. Once the number of active subscribers using a NAT IP address reaches this limit, that NAT IP address will not be allocated to new subscribers.



Important

In 19 and later releases, the number of users per NAT IP address can be configured dynamically for optimal utilization of NAT IP addresses.

Any new NAT IP allocated by VPN will take the configuration in the NAT pool. When a NAT IP is already allocated to Sessmgr, the change in configuration in NAT pool will not be applied to that NAT IP. This NAT IP will still use the value that was configured when it was allocated to Sessmgr by VPN. With NAT pool groups, each pool in a group must have the same number of users per NAT IP configured. If number of users per NAT IP configured in pools are different, though the configuration will still be allowed, the distribution of calls across pools in the group may not be even.

- **Port-chunk Thresholds:** The threshold is configured in percentage of total number of port-chunks. If the number of port-chunks already allocated from a given NAT IP address is less than the configured threshold limit of port-chunks, then the NAT IP address can be chosen for a new subscriber provided the “Maximum Users per NAT IP Address” is not reached. But if the number of chunks allocated is greater than or equal to the threshold limit of port-chunks, then the NAT IP address will not be chosen for a new subscriber. The remaining free port-chunks will be used for existing subscribers using the NAT IP address.

NAT IP Address Deallocation

Whenever a NAT IP address is deallocated, all the port-chunks associated with the subscriber are released back to the pool.

In case there is only one port-chunk associated with the subscriber:

- In case of many-to-one not-on-demand NAT IP pools, the last port-chunk is not released back to the pool even after NAT Binding Timer expires. Only when the call gets disconnected, the port-chunk is released along with the NAT IP address.
- In case of many-to-one on-demand NAT IP pools, when the last flow using the port-chunk gets cleared, the NAT Binding Timer is started. When the NAT Binding Timer expires, the port-chunk along with the NAT IP address is released back to the pool. NAT IP addresses can be forcibly released from SessMgr to VPNMgr for NAT pools using the **clear nat-ip** CLI command.
- In case of one-to-one on-demand NAT IP pools, when there are no active flows using a NAT IP address, the NAT Binding Timer is started. When the NAT Binding Timer expires, the NAT IP address gets deallocated.

NAT Port-chunk Allocation and Deallocation

This section describes the Port-chunk Allocation and Deallocation feature for many-to-one NAT.

NAT Port-chunk Allocation

Subscribers sharing a NAT IP address are allocated NAT ports in chunks. The ports in a port-chunk are always used for the subscriber to whom that port-chunk is allocated irrespective of the protocol.

Whenever a NAT IP address gets allocated to a subscriber, the first port-chunk gets allocated along with the NAT IP address. Thus, for not-on-demand pools, the first port-chunk gets allocated during call setup, and for on-demand pools during data flow.

A subscriber's TCP and UDP data traffic is NATed with ports chosen in a random fashion from the port-chunk allocated to that subscriber. For other protocol traffic, the first available port is allocated. When all the ports in a port-chunk are in use, a free port-chunk is requested for. A new port-chunk is only allocated if the "Maximum Port-chunks Per User" limit is not reached.

The port utilization data of subscribers is aggregated at the NAT pool level. The number of ports are grouped into buckets of size 8. There are 9 defined port buckets — [0-8], [9-16], [17-24], [25-32], [33-40], [41-48], [49-56], [57-64] and [≥ 65]. The first bucket [0-8] includes not-on-demand calls, that is, subscribers who are allocated a port chunk without using any ports at all will fall into the first bucket. The last bucket [≥ 65] includes all subscribers using greater than 64 ports. The maximum number of ports that were required by a subscriber at any point of time is recorded.



Important

In 19 and later releases, the port chunk size of an existing N:1 NAT pool can be changed dynamically without deleting or reconfiguring the pool.

When port chunk size is modified at pool level, it is possible that some NAT IP addresses will be already in use; for those active NAT IPs older value of port chunk size will be used. Only new NAT IPs being allocated from VPN will take the modified port chunk size. With NAT pool groups, each pool in a group must have the same port chunk size configured. If different chunk sizes are configured, though the configuration will still be allowed, the distribution of calls across pools in the group may not be even.

In release 19, the port chunk size is reduced to improve efficient usage of port chunks and NAT IP addresses allocated to a given Sessmgr. By increasing the number of users per NAT IP, the number of NAT IPs required to handle the calls in a given Sessmgr could come down. This will free NAT IP addresses and can be used by other Sessmgrs for allocation. The minimum port chunk size is reduced to 8 and the chunk size is configured in multiples of 8. The default chunk size will still be 32. If no chunk size is configured, then the chunk size is calculated by dividing the entire NAT port range by the number of users per NAT IP and then rounding off to the nearest multiple of 32.

NAT Port-chunk Deallocation

A port-chunk gets deallocated in the following cases:

- NAT Binding Timer expiry
- Subscriber session disconnect

NAT Binding Timer

When all flows using ports from a particular port-chunk get timed out/cleared, the port-chunk gets freed. When the last port of that port-chunk gets freed, the NAT Binding Timer starts counting. Before the NAT Binding Timer expires, if any new flows come up, ports are reallocated from the port-chunk, and the timer gets cancelled. The port-chunk cannot be deallocated as long as there are active flows using that port-chunk. But, if no new flows come and the NAT Binding Timer expires, the port-chunk gets deallocated.

In case of not-on-demand pools, the additional port-chunks that were allocated on demand will be deallocated based on the NAT binding timeout. However, the last port-chunk will not be deallocated even after the Binding Timer expires. This last port-chunk will only be deallocated when the NAT IP address is deallocated from the subscriber.

In case of on-demand pools, the port-chunks are deallocated based on the NAT binding timeout. When the last port-chunk gets freed, the NAT IP address also gets deallocated from the subscriber.

It is ensured that a port-chunk is associated with the subscriber as long as a valid NAT IP address is allocated to the subscriber.

Subscriber Session Disconnect

When a subscriber disconnects, all port-chunks associated with that subscriber are freed.

If the NAT Binding Timer has not expired, the port-chunks will not be usable immediately, only on NAT Binding Timer expiry will the port-chunks become available for new subscribers.

NAT IP Address/Port Allocation Failure

When a packet cannot be translated, the application can be notified by way of ICMP error messages, if configured. Translation failures may be due to no NAT IP address or port being available for translation.



Important

In the case of P-GW, NAT IP Address/Port Allocation Failure notification is not applicable.

TCP 2MSL Timer

NAT does port management only for many-to-one pools. Hence, The TCP 2MSL timer is only available for many-to-one NAT. It is necessary to ensure that a TCP NAT port in Time Wait state is not reused if there are other free ports available for the subscriber. If such a reuse happens, then there is a possibility that connections might get terminated by the server. To avoid such issues, whenever a many-to-one NAT TCP flow gets cleared, the NAT port goes to Time Wait state (2MSL started for that port). Once 2MSL timer expires, the NAT port becomes usable. The 2MSL timer is started for every TCP NAT port as soon as the TCP connection gets cleared. This ensures that a NAT TCP port gets reused only after expiry of the configured TCP 2MSL timer.

Consider a case where a single TCP flow is active in a port-chunk. When this connection gets cleared, the TCP NAT port goes to Time Wait state. Since this is the last flow of the port-chunk, the NAT Binding Timer also gets started.

Assume NAT Binding timer \geq TCP 2MSL timer. Once the 2MSL timer expires, the TCP port becomes usable. However, the NAT Binding Timer keeps counting, and on expiry, the port-chunk is released. In case the NAT Binding Timer $<$ TCP 2MSL Timer, on NAT Binding Timer expiry, the TCP port is forcefully moved to Free State (made usable) from Time Wait state and the port-chunk released.

Flow Mapping Timer

The Flow Mapping timer is a new timer implemented as an extension to the existing idle-timeout in ECS, and is supported only for TCP and UDP flows. This flow mapping applies only for NAT enabled calls.

The purpose of this timer is to hold the resources such as NAT IP, NAT port, and Private IP NPU flow associated with a 5-tuple ECS flow until Mapping timeout expiry. If the feature is disabled, the Flow mapping timeout will not get triggered for TCP/UDP idle timed out flows. The resources such as NAT mapping will be released with the 5-tuple flow itself.

NAT Binding Records

Whenever a NAT IP address or NAT port-chunk is allocated/deallocated to/from a subscriber, NAT Binding Records (NBR) can be generated. Generation of NBRs is configurable in the Firewall-and-NAT policy configuration.



Important

NAT Binding Records are now supported for NAT64.

NBRs are supported for both on-demand and not-on-demand NAT IP pools. For a one-to-one NAT IP pool, an NBR is generated whenever a NAT IP address is allocated/deallocated to/from a subscriber. For a many-to-one NAT IP pool, an NBR is generated when a port-chunk is allocated/deallocated to/from a subscriber for a NAT IP address. It is also possible to configure generation of NBRs only when a port-chunk is allocated, or deallocated, or in both cases.

NBRs can now hold both IPv4 and IPv6 addresses in case of an IPv4v6 subscriber. If the existing “ip subscriber-ip-address” is used for IPV4 or IPv4v6 call, IPv4 address will be generated and IPv6 address will be generated for IPv6 only call.

The following is the list of attributes that can be present in NBRs. You can configure a subset of these attributes or all of them to be logged in NBRs. If an attribute is not available, while logging records that field is populated with NULL.

- ip subscriber-ip-address: The private IP address.
- radius-calling-station-id: The IMSI of the mobile node.
- radius-fa-nas-identifier: A string that identifies PDSN. This field is optional if PDSN-NAS-IP address field is present.
- radius-fa-nas-ip-address:
- radius-user-name: NAI of the mobile node.
- sn-correlation-id: If available. The HA-Correlation-ID identifying the entire MIP session.
- sn-fa-correlation-id: If available. The PDSN-Correlation-ID as sent by the PDSN using the same format and length.
- sn-nat-binding-timer: Optional. The NAT Binding Timer assigned to the Realm.
- sn-nat-gmt-offset: Optional. The offset from GMT to correlate timestamps of records; GMT offset of the node generating this record. For example: -5.00, +5.30
- sn-nat-ip: The NAT IP address of mobile node.
- sn-nat-last-activity-time-gmt: The time the last flow in a specific NAT set of flows was seen in GMT time.
- sn-nat-port-block-end: The NAT Port Block End of the mobile node.
- sn-nat-port-block-start: The NAT Port Block Start of the mobile node.

- sn-nat-port-chunk-alloc-dealloc-flag: 1: allocate; 0: deallocate
- sn-nat-port-chunk-alloc-time-gmt: The NAT Port Chunk Allocation Timestamp (Sample time format: 03/11/2009 10:38:35)
- sn-nat-port-chunk-dealloc-time-gmt: The NAT Port Chunk Deallocation Timestamp (Sample time format: 03/11/2009 10:38:35)
- sn-nat-realm-name: Optional. The name of the locally configured NAT Realm.
- sn-nat-subscribers-per-ip-address: Optional. NAT Multiplier assigned to the Realm.
- subscriber-ipv4-address: The subscriber IPv4 address in the NBR.
- subscriber-ipv6-address: The subscriber IPv6 prefix address in the NBR.
- bearer 3gpp charging-id: The charging ID for the PDN Session.
- bearer 3gpp sgsn-address: The S-GW/SGSN address.
- bearer ggsn-address: The P-GW/GGSN address.
- bearer 3gpp imsi: The IMSI value of the subscriber.

**Important**

The NBR attributes: sn-correlation-id, sn-fa-correlation-id, radius-fa-nas-ip-address, radius-fa-nas-identifier are not applicable for P-GW and GGSN.

Bulk Statistics Support

Bulk statistics for NBRs are supported in the ECS schema. These bulk statistics are collected when NBRs are generated for IP/Port chunk allocations/deallocations.

- total-nbrs-generated
- nbrs-for-port-chunk-alloc
- nbrs-for-port-chunk-release

NAT Binding Updates

Whenever a NAT IP address or NAT port-chunk is allocated/deallocated to/from a subscriber, to update NAT binding information for that subscriber in the AAA, a NAT Binding Update (NBU) can be sent to the AAA server.

**Important**

NAT Binding Updates are not supported for NAT64.

**Important**

P-GW and GGSN do not support the NBU feature.

Since port-chunk allocation/deallocation happens on a per-call basis, this ensures that AAA messaging is reduced to a great extent. NBUs are sent to the AAA server in accounting-interim messages. To send or not to send NBUs to the AAA server is configurable in the NAT IP pool configuration.

NBUs are supported for both one-to-one and many-to-one NAT IP pools.

An NBU contains the following attributes:

- Alloc-Flag
- Binding-Timer
- Correlation-Id
- Loading-Factor
- NAT-IP-Address
- NAT-Port-Block-End: In the case of one-to-one NAT, the value is 65535
- NAT-Port-Block-Start: In the case of one-to-one NAT, the value is 1

CoA NAT Query

If the NAT binding information is not available at the AAA, the AAA server can query the chassis for the information. This query uses the Change of Authorization (CoA) format, wherein the AAA sends a one-to-one NAT IP address as a query, and in the CoA query response the NBU is obtained if available at the time of query.



Important CoA NAT Query is not supported for NAT64.



Important CoA query for NAT binding information is only supported for one-to-one NAT.

The CoA query request must contain the following attributes:

- Event-Timestamp
- NAS-IP-Address
- SN1-NAT-IP-Address

For SN1-NAT-IP-Address, supported VSA-Type values 0 and 1.

For a successful query, the CoA ACK response contains the following attributes:

- Acct-Session-Id
- Correlation-Id
- Framed-IP-Address
- NAT-IP-Address
- NAT-Port-Block-End
- NAT-Port-Block-Start
- User-Name



Important For more information on the AVPs/VSAs, if you are using StarOS 12.3 or an earlier release, please refer to the *AAA and GTPP Interface Administration and Reference*. If you are using StarOS 14.0 or a later release, refer to the *AAA Interface Administration and Reference*.

Firewall-and-NAT Policy

A Firewall-and-NAT policy contains a set of access ruledefs with priorities and actions, and the NAT configurations. On a system, multiple such policies can be configured, however at any point of time only one

policy is associated to a subscriber. Firewall-and-NAT policies are configured in the CLI Firewall-and-NAT Policy Configuration Mode.

**Important**

In release 8.x, NAT for CDMA and early UMTS releases used rulebase-based configurations, whereas in later UMTS releases NAT used policy-based configurations. In 9.0 and later releases, NAT for UMTS and CDMA releases both use policy-based configurations. For more information, please contact your local service representative.

**Important**

In a Firewall-and-NAT policy, a maximum of twenty NAT IP pools/NAT IP pool groups can be configured. At any time a subscriber can be associated with a maximum of three different NAT IP pools/NAT IP pool groups and can have NATed flows on three different NAT IP addresses at the same time.

New NAT IP pools/NAT IP pool groups cannot be added to a policy if the maximum allowed is already configured in it. However, a pool/pool group can be removed and then a new one added. When a pool/pool group is removed and a new one added, the pool/pool group that was removed will stay associated with the subscriber as long as the subscriber has active flows using that pool/pool group. If the subscriber is already associated with three NAT IP pools (maximum allowed), any new flows from that subscriber for the newly added pool will be dropped. A deleted pool is disassociated from the subscriber on termination of all flows from that subscriber using that pool. The new pool/pool group is associated with the subscriber only when the subscriber sends a packet to the newly added pool.

In the Firewall-and-NAT policy configuration, the NAT44/NAT64 policy must be enabled. Once NAT is enabled for a subscriber, the NAT IP address to be used is chosen from the NAT IP pools/NAT IP pool groups specified in matching access rules configured in the Firewall-and-NAT policy.

The Firewall-and-NAT policy used for a subscriber can be changed either from the command line interface, or through dynamic update of policy name in Diameter and RADIUS messages. In both the cases, NAT status on the active call remains unchanged.

The Firewall-and-NAT policy to be used for a subscriber can be configured in:

- ECS Rulebase: The default Firewall-and-NAT policy configured in the ECS rulebase has the least priority. If there is no policy configured in the APN/subscriber template, and/or no policy to use is received from the AAA/OCS, only then the default policy configured in the ECS rulebase is used.
- APN/Subscriber Template: The Firewall-and-NAT policy configured in the APN/subscriber template overrides the default policy configured in the ECS rulebase. To use the default policy configured in the ECS rulebase, in the APN/subscriber configuration, the command to use the default rulebase policy must be configured.
- AAA/OCS: The Firewall-and-NAT policy to be used can come from the AAA server or the OCS. If the policy comes from the AAA/OCS, it will override the policy configured in the APN/subscriber template and/or the ECS rulebase.

**Important**

The Firewall-and-NAT policy received from the AAA and OCS have the same priority. Whichever comes latest, either from AAA/OCS, is applied.

The Firewall-and-NAT policy to use can also be received from RADIUS during authentication.

Disabling NAT Policy



Important By default, NAT processing for subscribers is disabled.

NAT processing for subscribers is disabled in the following cases:

- If the AAA/OCS sends the SN-Firewall-Policy AVP with the string “disable”, the locally configured Firewall-and-NAT policy does not get applied.
- If the SN-Firewall-Policy AVP is received with the string “NULL”, the existing Firewall-and-NAT policy will continue.
- If the SN-Firewall-Policy AVP is received with a name that is not configured locally, the subscriber session is terminated.

Updating Firewall-and-NAT Policy in Mid-session

The Firewall-and-NAT policy can be updated mid-session provided the policy was enabled during call setup. Firewall-and-NAT policy can also be updated during mid-session rulebase update if the Firewall-and-NAT policy was previously assigned through rulebase.



Important When the firewall AVP contains “disable” during mid-session firewall policy change, there will be no action taken as the Firewall-and-NAT policy cannot be disabled dynamically. The policy currently applied will continue.



Important For all NAT/Firewall-enabled subscribers, when the Firewall-and-NAT policy is deleted, the call is dropped.

In a Firewall-and-NAT policy, you can change the NAT enabled/disabled status at any time. However, the updated NAT status will only be applied to new calls, active calls using that Firewall-and-NAT policy will remain unaffected.

Target-based NAT Configuration

A NAT IP pool can be selected based on the L3/L4 characteristics of a subscriber’s flows. NAT can be configured such that all subscriber traffic coming towards specific public IP address(es) always selects a specific NAT IP pool based on the L3/L4 traffic characteristics.



Important A subscriber can be allocated only one NAT IP address per NAT IP pool/NAT IP pool group from a maximum of three NAT IP pools/NAT IP pool groups. Hence, at anytime, there can only be a maximum of three NAT IP addresses allocated to a subscriber.

This association is done with the help of access ruledefs configured in the Firewall-and-NAT policy. The NAT IP pool/NAT IP address to be used for a subscriber flow is decided during rule match. When packets match an access ruledef, NAT is applied using the NAT IP address allocated to the subscriber from the NAT IP pool/NAT IP pool group configured in that access ruledef.

If no NAT IP pool/NAT IP pool group name is configured in the access ruledef matching the packet, and if there is a NAT IP pool/NAT IP pool group configured for “no ruledef matches”, a NAT IP address from the NAT IP pool/NAT IP pool group configured for “no ruledef matches” is allocated to the flow.

If no NAT IP pool/NAT IP pool group is configured for “no ruledef matches” and if there is a default NAT IP pool/NAT IP pool group configured in the rulebase, a NAT IP address from this default NAT IP pool/NAT IP pool group is allocated to the flow.

If a NAT IP pool/NAT IP pool group is not configured in any of the above cases, no NAT will be performed for the flow. Or, if bypass NAT is configured in a matched access rule or for “no ruledef matches” then NAT will not be applied even if the default NAT IP pool/NAT IP pool group is configured. The order of priority is:

1. Bypass NAT
2. NAT IP pool/NAT IP pool group in ruledef
3. NAT IP pool/NAT IP pool group for “no-ruledef-matches”
4. Default NAT IP pool/NAT IP pool group

When a new NAT IP pool/NAT IP pool group is added to a Firewall-and-NAT policy, it is associated with the active subscriber (call) only if that call is associated with less than three (maximum limit) NAT IP pools/NAT IP pool groups. If the subscriber is already associated with three NAT IP pools/NAT IP pool groups, any new flows referring to the newly added NAT IP pool/NAT IP pool group will get dropped. The newly added NAT IP pool/NAT IP pool group is associated to a call only when one of the previously associated NAT IP pools/NAT IP pool groups is freed from the call.

NAT Application Level Gateway

Some network applications exchange IP/port information of the host endpoints as part of the packet payload. This information is used to create new flows, by server or client.

As part of NAT ALGs, the IP/port information is extracted from the payload, and the flows are allowed dynamically (through pinholes). IP and port translations are done accordingly. However, the sender application may not be aware of these translations since these are transparent, so they insert the private IP or port in the payload as usual. For example, FTP NAT ALG interprets “PORT” and “PASV reply” messages, and NAT translates the same in the payload so that FTP happens transparently through NAT. This payload-level translation is handled by the NAT ALG module.

The NAT module will have multiple NAT ALGs for each individual application or protocol.

Supported NAT ALGs

NAT ALGs are supported only for the following protocols:

- H323
- File Transfer Protocol (FTP)
- Point-to-Point Tunneling Protocol (PPTP): If PPTP ALG is enabled, NAT is supported for GRE flows that are generated by PPTP.
- Real Time Streaming Protocol (RTSP)
- Session Initiation Protocol (SIP)

- Trivial File Transfer Protocol (TFTP)

For NAT ALG processing, in the rulebase, routing rules must be configured to route packets to the corresponding analyzers.

Session recovery is supported for ALG. Only one contact pinhole, and only one connected call and its associated media pinholes will be recovered for a subscriber. Any subscriptions, ongoing transactions, or unconnected calls will not be recovered. SIP ALG recovery data will be check-pointed using the variable length micro checkpointing mechanism.

SIP ALG is made compatible with user-to-user authentication and processing 4xx responses as described in RFC 3261 (SIP - Session Initiation Protocol).

SIP and H323 ALGs support multiple IPs per NAT realm and other ALGs (FTP, PPTP, TFTP, RTSP) do not support multiple IPs per NAT realm.

H323 ALG Support

H323 ALG is supported to traverse NAT by inspecting and altering information contained in existing H323 messages as they pass through the NAT. It can alter address and port information in registration, call signaling and automatically open pinholes in the NAT to allow media flow.

H323 ALG performs the following functions:

- Communicates with the core for binding management
- Communicates with NAT for signaling messages
- Uses H323 stack for parsing and encoding the H323 messages
- Performs protocol specific processing if required

The following supplementary services are currently supported in H323 ALG:

- Call Transfer: The Call Transfer supplementary service enables the served user (User A) to transform an existing call with a User B (primary call) into a new call between current User B and a new User C (transferred-to user) selected by served user A.
- Call Hold: The Call Hold supplementary service allows the served user, which may be the originally calling or the called user, to interrupt communications on an existing call and then subsequently, if desired, re-establish (i.e. retrieve) communications with the held user.
- Call Diversion: Call Diversion supplementary service permits a served user to have incoming calls addressed to the served user's number redirected to another number; on busy service, it enables a served user to have calls redirected to another endpoint; on No Answer, it enables a served user to have calls addressed to the served endpoint's number and redirected to another endpoint if the connection is not established within a defined period of time.
- Call Waiting: The Call Waiting supplementary service permits a busy user to be informed of an incoming call while being engaged with one or more other calls.
- Call Offering: The Call Offering supplementary service on request from the calling user, enables a call to be offered to a busy user and to wait for that called user to accept the call, after the necessary resources have become available.

NAT Aware H323 Clients

An application layer gateway, at the Firewall/NAT, examines all the H323 packets and modifies the packet such that all the private addresses are replaced by public addresses. It also opens all the pinholes required for successful call establishment. A NAT aware endpoint establishes end-to-end media session through FW/NAT without the need of ALG. Any TCP connection or UDP packet sent from the internal network through the firewall opens a pinhole dynamically in the firewall. This pinhole allows incoming messages to be sent from the destination of the TCP connection or the UDP packet. The pinhole stays open as long as the network sends information through the pinhole to the same destination.

If an end point supports NAT traversal, H323 ALG disables itself so that end point directly opens required pinhole and establishes media path between them. The ALG will not manage any pinhole for media traversal across Firewall/NAT for NAT aware clients. By default, the ALG will bypass all the clients that support H460.18/19 and H460.23/24.

Accelerated ECS Feature Support

Accelerated-ECS (A-ECS) feature speeds up the processing of certain types of flows such that packet-actions and charging applicable to packets from those flows is done in a fast manner. The throughput in terms of PPS (Packets Processed per Second) is improved by caching rule matching results for a flow for selected flows so as to not incur the lookup penalty for a large number of packets in that flow. The A-ECS path is capable of performing a full range of basic functions including handling charging, modification of packet headers and incrementing various counters. Accelerated ECS identifies packets that need only a small amount of processing, and performs only those necessary tasks on these packets. Only those packets that do not require DPI are allowed to enter the Accelerated path.

Layer-3/Layer-4 NAT interworks with A-ECS, so that A-ECS will capture a larger chunk of traffic at various operators that use L3/L4 NAT. This basically involves separating out the NAT and SFW functionalities, and within that, separate out L3/L4 NAT from ALG-NAT. Once that is done, the accelerated-path is modified to allow L3/L4 NAT flows, and not SFW or ALG flows.

For more information on the Accelerated-ECS, refer to the *ECS Administration Guide*.

EDRs and UDRs

This section describes the NAT-specific attributes supported in EDRs and UDRs.

EDRs

The following NAT-specific attributes are supported in regular EDRs:

- sn-nat-subscribers-per-ip-address: Subscriber(s) per NAT IP address
- sn-subscriber-nat-flow-ip: NAT IP address of NAT-enabled subscribers
- sn-subscriber-nat-flow-port: NAT port number of NAT-enabled subscribers

UDRs

The following NAT-specific attribute is supported in regular UDRs:

sn-subscriber-nat-flow-ip: NAT IP addresses that are being used by NAT-enabled subscribers. The NAT IP addresses assigned from each of the associated pool for the call are logged. A space is used as a separator between individual IP addresses.

Bulk Statistics

The NAT realms are configured in a context and statistics are stored per context per realm. These statistic variables, both cumulative and snapshot, are available in the nat-realm schema.

Bulkstats are only generated for the first 100 NAT IP pools from an alphabetical list of all NAT IP pools, which is based on the context name and pool name. Therefore, to generate bulkstats for a specific NAT IP pool it must be named such that it gets selected in the first 100 bulkstats.

The following are cumulative statistics that can be part of NAT bulkstats:

- vppname: Context name
- realmname: Realm name
- nat-bind-updates: Total interim AAA NBU sent.
This is available only in StarOS 12.3 and earlier releases.
- nat-rlm-bind-updates: Total interim AAA NBU sent.
This is available only in StarOS 14.0 and later releases.
- nat-rlm-bytes-tx: Total number of NAT44 and NAT64 bytes transferred by realm (uplink + downlink).
This is available only in StarOS 12.3 and earlier releases.
- nat-rlm-bytes-txferred: Total number of NAT44 and NAT64 bytes transferred by realm (uplink + downlink).
This is available only in StarOS 14.0 and later releases.
- nat-rlm-bytes-nat44-tx: Total number of NAT44 bytes transferred by realm.
- nat-rlm-bytes-nat64-tx: Total number of NAT64 bytes transferred by realm.
- nat-rlm-flows: Total number of NAT44 and NAT64 flows used by the realm.
This is available only in StarOS 12.3 and earlier releases.
- nat-rlm-ip-flows: Total number of NAT44 and NAT64 flows used by the realm.
This is available only in StarOS 14.0 and later releases.
- nat-rlm-nat44-flows: Total number of NAT44 flows processed by realm.
- nat-rlm-nat64-flows: Total number of NAT64 flows processed by realm.
- nat-rlm-ip-denied: Total number of NAT44 and NAT64 flows denied NAT IP address.
- nat-rlm-ip-denied-nat44: Total number of NAT44 flows denied IP.
- nat-rlm-ip-denied-nat64: Total number of NAT64 flows denied IP.
- nat-rlm-port-denied: Total number of NAT44 and NAT64 flows denied ports.
- nat-rlm-port-denied-nat44: Total number of NAT44 flows denied ports.

- nat-rlm-port-denied-nat64: Total number of NAT64 flows denied ports.
- nat-rlm-max-port-chunk-subs: Total number of subscribers who used maximum number of port chunks.
- nat-rlm-max-port-chunk-used: Maximum port chunks used.
- nat-rlm-memory-denied: Total number of NAT44 and NAT64 flows denied memory.
- nat-rlm-memory-denied-nat44: Total number of NAT44 flows denied memory.
- nat-rlm-memory-denied-nat64: Total number of NAT64 flows denied memory.

The following are snapshot statistics that can be part of NAT bulkstats:

- vpnname: Context name
- realmname: Realm name
- nat-rlm-ttl-ips: Total number of NAT public IP addresses, per context per NAT realm. Is a static value.
- nat-rlm-ips-in-use: Total number of NAT IP addresses currently in use, per context per NAT realm.
- nat-rlm-current-users: Total number of subscribers currently using the NAT realm.
- nat-rlm-ttl-port-chunks: Total number port-chunks, per context per NAT realm. Is a static value.
- nat-rlm-chunks-in-use: Total number of port-chunks currently in use, per context per NAT realm.
- nat-rlm-max-cur-port-chunk-subs: Current number of subscribers using maximum number of port chunks.
- nat-rlm-max-cur-port-chunk-used: Maximum port chunks used by active subscribers.
- nat-rlm-port-chunk-size: Size of the port chunk in the NAT realm.
- nat-rlm-port-chunk-average-usage-tcp: Average TCP port usage in the allocated TCP ports, i.e. out of allocated TCP ports how many got used. Not percentage value.
- nat-rlm-port-chunk-average-usage-udp: Average UDP port usage in the allocated UDP ports, i.e. out of allocated UDP ports how many got used. Not percentage value.
- nat-rlm-port-chunk-average-usage-others: Average other (ICMP or GRE) port usage in the allocated other ports, i.e. out of allocated 'other' ports how many got used. Not percentage value.

Alarms

Alert threshold values can be specified to generate alarms for NAT IP pools. To specify realm-specific threshold limits (pool-used, pool-free, pool-release, and pool-hold) "alert-threshold" NAT IP pool parameter can be used, or it can also be specified across context. These thresholds can be specified to any number of NAT IP pools.

In case of many-to-one NAT, it is possible to specify port-chunks usage threshold per NAT IP pool. This threshold value is applicable to all many-to-one NAT IP pools across the system. However, note that alarms are only generated for the first 100 many-to-one NAT IP pools from an alphabetical list of all NAT IP pools.

Session Recovery and ICSR

In session recovery, as part of the Private IP assigned to the subscriber:

- The public IP address used for the subscriber is recovered. The NAT IP address being used by the subscriber can be on-demand or not-on-demand. In case of many-to-one NAT, the port-chunks associated with the NAT IP address for the subscriber needs to be checkpointed as well.
- In case Bypass NAT feature is used, then the private IP flow needs to be recovered.

To be recovered the NAT IP addresses need to be checkpointed. The checkpointing can be:

- Full Checkpoint
- Micro Checkpoint

To recover the bypass NAT flow, the bypass flow needs to be checkpointed. The checkpointing of Bypass NAT flow can be:

- Full Checkpoint
- Micro Checkpoint

In case of not-on-demand, the NAT IP address being used by the subscriber is known after call setup. This gets checkpointed as part of the normal full checkpoint. In case of on-demand NAT, the NAT IP address being used by the subscriber is known only in the data-path. This will be checkpointed as part of micro checkpoint.

In case of many-to-one NAT, the port-chunks being used will always be checkpointed as part of micro checkpoint.

In case of bypass NAT flow, in most cases the flow gets checkpointed as part of micro checkpoint.

Any information that is checkpointed as part of full checkpoint is always recovered. Data checkpointed through micro checkpoint cannot be guaranteed to be recovered. The timing of switchover plays a role for recovery of data done through micro checkpoint. If failover happens after micro checkpoint is completed, then the micro checkpointed data will get recovered. If failover happens during micro checkpoint, then the data recovered will be the one obtained from full checkpoint.

Once NAT IP/and Port-Chunks/Bypass NAT flow are recovered, the following holds good:

- One-to-one NAT: Since NAT IP address being used for one-to-one NAT is recovered, on-going flows will be recovered as part of Firewall Flow Recovery algorithm as one-to-one NAT does not change the port.
- Many-to-one NAT: On-going flows will not be recovered as the port numbers being used for flows across chassis peers/SessMgr peers are not preserved.

It is now possible to enable/disable the checkpointing of NATed flows and control the type of flows to be checkpointed based on criteria. Check pointing is done only for TCP and UDP flows.

Many-to-one NAT flow recovery is supported for ICSR.

- Bypass NAT Flow: On-going flows will be recovered as part of Firewall Flow Recovery algorithm.

All of the above items is applicable for ICSR as well. SIP ALG also supports ICSR and is applicable only to UDP flows.

In Firewall-and-NAT policy, checkpointing and ICSR recovery for basic NAT, SIP and H323 flows can be configured. A maximum of 100 basic flows can be checkpointed.


Category	Event	Impacted	Details
One-to-One NAT	Session	No	Session recovered.
	New Traffic	No	NAT will be applied.
	Ongoing Traffic	Yes	Cannot differentiate between ongoing traffic and unsolicited traffic. A rule-match is done and if allowed, NAT will be applied accordingly on the packet.
	Unsolicited Traffic (downlink packets)	Yes	Cannot differentiate between ongoing traffic and unsolicited traffic. Translation will be done and packet action taken based on the rule-match.
Many-to-One NAT	Session	No	Session recovered.
	New Traffic	No	NAT will be applied.
	Ongoing Traffic — TCP	Yes	Packet will be dropped.
	Ongoing Traffic — UDP	Yes and No	If it is downlink packet, it will be dropped. If it is uplink packet, NAT will be applied with a new port.
	Ongoing Traffic — ICMP	Yes and No	If it is downlink packet, it will be dropped. If it is uplink packet, NAT will be applied with a new port.
	Unsolicited Traffic (downlink packets)	No	Packet will be dropped.
Bypass NAT	Session	No	Session recovered.
	New Traffic	No	Traffic will be NAT bypassed.
	Ongoing Traffic	No	Traffic will be NAT bypassed.
	Unsolicited Traffic (downlink packets)	No	Traffic will be NAT bypassed.

For more information, see the *Session Recovery* and *Interchassis Session Recovery* chapters in the *System Administration Guide*.

NAT64 Overview

Stateful NAT64 is a mechanism for translating IPv6 packets to IPv4 packets and vice-versa. The IPv4 address of IPv4 server/host in an IPv4 network is obtained to and from IPv6 addresses by using the configured stateful prefix. The IPv6 addresses of IPv6 hosts are translated to and from IPv4 addresses by installing mappings in the usual NAT manner. The following figure illustrates the working of NAT64 with DNS64.

Figure 1: NAT64 Mechanism



NAT64 is applied on traffic based on the rule match (Destination based NATing). If NAT64 has to be applied, then the NAT64 will translate and forward them as IPv4 packets through the IPv4 network to the IPv4 receiver. The reverse takes place for packets generated by hosts connected to the IPv4 network for an IPv6 receiver. If NAT64 is not applied on the IPv6 packet, then the IPv6 packet will not be translated and sent as is (NAT bypassed) and will be routed within the IPv6 network to the destination.

NAT64 will not be applied for packets whose destination IP address does not match a pre-defined prefix. NAT64 will be applied only for packets whose destination IP address matches a pre-defined prefix. The pre-defined prefix is configurable and it is a single prefix.

To summarize, the IPv4-over-IPv6 solution or the 464XLAT feature is one the technique used to support IPv4 service extension and IPv6 deployment. 464XLAT uses the IPv4/IPv6 standardized translation (used in RFC6145 and RFC6146). It does not require DNS64 (RFC6147) because an IPv4 host may simply send IPv4 packets, including packets to an IPv4 DNS server, that will be translated to IPv6 on the customer-side translator (CLAT) and back to IPv4 on the provider-side translator (PLAT). 464XLAT networks may use DNS64 (RFC6147) to enable single stateful translation (RFC6146) instead of 464XLAT double translation where possible. It facilitates the IPv6 transition by making IPv4 services reachable across IPv6-only networks and provides IPv6 and IPv4 connectivity to single-stack IPv4 or IPv6 servers and peers.

In Release 21.2, this feature is implemented on the P-GW, which uses the PLAT functionality with the NAT64 solution.

The P-GW identifies and provides proper NAT64 for user data plane traffic destined for IPv4 networks. The destination IPv6 address will contain the predefined IPv6 address prefix given by DNS64. The P-GW configured with the same IPv6 prefix will only send those data packet to the NAT64 process and forward it to an IPv4 network for transport to end destination.

NAT64 Translation

For NAT64, Network address translation and Protocol translation are done on the packets. The uplink IPv6 packets that are destined to hosts in the IPv4 network must be protocol translated to IPv4 packets and forwarded. The downlink IPv4 packets destined to hosts in IPv6 network must be protocol translated to IPv6 packets and then forwarded.

The Network address translation is done using the following ways:

- **One-to-One NAT:** In the case of 1:1 NAT, the subscriber IPv6 address is uniquely mapped to a given NAT IPv4 address. Port translation is not done as the NAT IP address is associated with a single subscriber and not shared by many users.

One-to-One NAT IP allocated to a subscriber can be simultaneously used for NATing IPv4 traffic and IPv6 traffic from a given subscriber. When downlink packets are received, firstly the NAT64 binding lookup is performed for NAT64 translation. If lookup is not successful, then the packet will be NAT44 translated.

- **Many-to-One NAT:** In the case of N:1 NAT, the subscriber IPv6 address and source port is mapped to a given NAT IPv4 address and NAT port. Port translation must be done as the same NAT IPv4 address is shared by multiple users. Hence, the L4 ports must be translated to differentiate the connections originating from multiple users sharing the same NAT IPv4 address.

Limitations for One-to-One NAT64

This section lists the limitations for One-to-One NAT64.

- In the case of One-to-One NAT, a given destination can be associated with only one prefix at any point of time as maintained in the destination prefix list. If the same destination has to be associated with multiple prefixes, then such packets will be dropped.
- Any downlink traffic received on One-to-One NAT IP will always be translated to the same 128-bit IPv6 address (though interface IDs can actually be different).
- One-to-One NAT IP status is lost after recovery. The NAT IP that was previously used for NAT44 or NAT64 is not recovered. Based on the first packet that is received after call recovery and the PDN type, the IP will be used for NATing IPv4 or IPv6 traffic.

Protocol Translation

This section describes the Uplink and Downlink Packet translation.

- **Uplink Packet Translation:** The uplink packets are translated from IPv6 to IPv4. The IP headers in the packet will be translated. The existing NAT APIs are enhanced to perform Protocol translation. Along with the NAT mapping, the prefix/suffix to be used for translation will also be passed. In case of fragmented packets, the packets need to be reassembled and then translated. The uplink packet translation includes:
 - **IPv6 to IPv4 Header Translation:** The original IPv6 header on the packet is removed and replaced by an IPv4 header.
 - **ICMPv6 to ICMPv4 Header Translation:** The original ICMPv6 header on the packet is removed and replaced by an ICMPv4 header.
 - **Packet Translation**
- **Downlink Packet Translation:** The downlink packets need to be translated from IPv4 to IPv6. The existing NAT APIs are to be enhanced to perform Protocol translation. Along with the NAT mapping, the prefix/suffix to be used for translation will also be passed. In case of fragmented packets, the packets need to be reassembled and then translated. The downlink packet translation includes:
 - **IPv4 to IPv6 Header Translation:** The original IPv4 header on the packet is removed and replaced by an IPv6 header.
 - **ICMPv4 to ICMPv6 Header Translation:** The original ICMPv4 header on the packet is removed and replaced by an ICMPv6 header.

NAT64 ALGs Support

NAT64 ALGs support the following protocols:

- File Transfer Protocol (FTP)
- Point-to-Point Tunneling Protocol (PPTP)
- Real Time Streaming Protocol (RTSP)
- Session Initiation Protocol (SIP)
- Trivial File Transfer Protocol (TFTP)

ICMP Host Unreachable

In earlier releases, the NAT44 and NAT64 features could not forward downlink-initiated flows on ASR 5500 because of the unknown public IP address, port, private IP address and port binding, and as a result the packets were getting dropped. The Internet server and other nodes that forward the packets from the Internet server to the ASR 5500 are unaware of this drop, and maintain the flow-related parameters for the dropped packet.

To resolve the condition of packets getting dropped, NAT44/NAT64 now sends ICMP Host Unreachable for all downlink packets that get dropped. In the case of Many-to-Many NAT, there are chances of downlink packets getting dropped when there is no existing flow. In Many-to-Many NAT, downlink packets will be considered as unsolicited under the following conditions:

- No NAT Binding exists.
- Binding exists but there is no active 5 tuple flow.

In case of One-to-One NAT, downlink packets will be considered as unsolicited under the condition that there is no 5-tuple flow. With ICMP-HU feature enabled, NAT sends ICMP-HU after dropping the unsolicited packets.

Port Control Protocol Support

The Port Control Protocol (PCP) feature provides a mechanism to control how incoming packets are forwarded by upstream devices such as Network Address Translation IPv4/IPv4 (NAT44) and IPv4 firewall devices, and to reduce application keepalive traffic.



Important

The PCP feature is customer specific. For more information contact your Cisco account representative.



Important


PCP is a licensed Cisco feature. Contact your Cisco account representative for more information. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

The PCP server is supported on ASR 5500 chassis running in-line services such as NAT44 and Stateful Firewall(s) individually or in collocated configurations. PCP supports the following functions:

- A host to control how incoming packets are forwarded by upstream devices such as Network Address Translation (NAT44) and Stateful Firewall (IPv4).
- A host to reduce the application keepalive messages.
- A host to operate a server for a long duration (e.g. webcam) or a short duration (e.g. while playing a game or on a phone call) when behind a NAT device, including a CGN operated by an Internet service provider.
- Applications to create mappings from an external IP address and port to an internal (target) IP address and port. These mappings are required for successful inbound communications destined to machines located behind a NAT or Firewall.
- Applications to create mappings in NAT and Firewall, and reducing the incentive to deploy ALGs in NAT and Firewalls.

The following figure shows NAT44 and PCP Server on ASR 5500.

Figure 2: NAT44 and PCP Server



The PCP service has to be associated with a PCP server IP address. The PCP server IP address is picked from the destination context associated with the subscriber. Only, if such an IP address is available and the status is up, the PCP service will listen to PCP requests on that IP address. The PCP service will be bound only to an IPv4 address and listens on UDP port (5351 (default port) or can be configured).

In case of system failure, the PCP service recovers along with subscriber's PCP enabled status. In case of stand-alone recovery and ICSR, only the subscriber PCP enabled status will be check-pointed.

PCP supports interworking with the following existing NAT ALGs:

- FTP
- RTSP
- SIP

Bulk Statistics Support

Bulk statistics reporting for the PCP feature is supported.

For the PCP feature the following bulk statistics are available in the ECS schema:

- total-pcp-svc-req
- total-pcp-svc-rsp
- total-pcp-svc-unknown-rsp
- total-pcp-svc-invalid-rsp
- total-pcp-svc-map-req

- total-pcp-svc-map-valid-req
- total-pcp-svc-map-invalid-req
- total-pcp-svc-map-rsp
- total-pcp-svc-map-rsp-success
- total-pcp-svc-peer-rsp-error
- total-pcp-svc-peer-req
- total-pcp-svc-peer-valid-req
- total-pcp-svc-peer-invalid-req
- total-pcp-svc-peer-rsp
- total-pcp-svc-peer-rsp-success
- total-pcp-svc-peer-rsp-error
- total-pcp-svc-announce-req
- total-pcp-svc-announce-valid-req
- total-pcp-svc-announce-invalid-req
- total-pcp-svc-announce-rsp
- total-pcp-svc-announce-rsp-success
- total-pcp-svc-announce-rsp-error
- total-pcp-svc-subscribers
- current-pcp-svc-subscribers

Logging Support

NAT supports logging of various messages on screen if logging is enabled for NAT. These logs provide detailed messages at various levels, like critical, error, warning, and debug. NAT attack logs also provide information on the source IP address, destination IP address, protocol, or attack type for any packet dropped due to an attack. These logs are also sent to a syslog server if configured in the system.

Enhanced Syslog Reporting

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	All Products
------------------------------------------	--------------

Applicable Platform(s)	<ul style="list-style-type: none"> • ASR 5500 • VPC - Di • VPC - Si
Feature Default	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>ASR 5500 System Administration Guide</i> • <i>NAT Administration Guide</i> • <i>PSF Administration Guide</i>

Revision History

Revision Details	Release
With this feature, the port information of the dropped packet is included in the logs.	21.3.1
First introduced.	Pre 21.2

Feature Description

Firewall and NAT attack logs provide information on the source IP address, destination IP address, protocol, or attack type for any packet dropped due to an attack. Prior to this release, when an attack happened, the logs did not carry any information about the ports.

With this feature, the port information of the dropped packet is included in the logs. The port information such as source port and destination port are important while configuring access rules to allow or block certain traffic.

Following are some important points to be considered:

- Typically, attack logs are at INFO/DEBUG level. At this level, there are too many logs generated even for normal traffic. Hence, to segregate the attack logs, the attack logs are moved to the WARNING level.
- Firewall and NAT attack logs are moved to WARNING level from Info/Debug level. The source port and destination port are logged as part of Firewall/NAT attack logs.
- Both IPv4 and IPv6 traffic is supported.
- The source port and destination port are valid for TCP/UDP protocols. However, for other protocols, the ports are logged as zero.

Previous Behavior: Earlier, the attack logs did not carry any port information and the logs were of the type Info/Debug.

New Behavior: With this feature, firewall and NAT attack log levels have been changed to WARNING from INFO/DEBUG for event IDs 96188, 96995, 96186, 96185, 96159, and 96203. Source port and destination port information are now displayed in the attack logs.

Impact on Customer: The attack logs are per packet logs seen at WARNING level. If you enable WARNING and above logs for Firewall (NAT) facility, and when there is an attack, log rate is very high.

Limitations

Following are the limitations of this feature:

- The attack logs are per packet logs and when an attack is in progress, log generation rate is very high.
- Under extreme attack conditions, evlogd CPU might go up.
- Event IDs for the attack logs:
 - firewall 96188 warning
 - firewall 96995 warning
 - firewall 96186 warning
 - firewall 96203 warning
 - firewall 96159 warning
 - firewall 96185 warning
- When there are too many logs generated under attack conditions, the following event IDs must be disabled:
 - eventid 96188: Disables Firewall Attack log generation
 - eventid 96186: Disables Port Scan Attack log generation
 - eventid 96995: Disables NAT Attack log generation
 - eventid 96203: Disables logging for TCP reset message threshold breach
 - eventid 96159: Disables logging for packets denied by rule
 - eventid 96185: Disables logging for ICMP unreachable message threshold breach

Configuring Logging Event ID

When there are too many logs generated under attack conditions, use the following command to disable the event IDs:

- To disable firewall attack log generation:

```
[local]asr5500(config)# logging disable eventid 96188
```
- To disable port scan log generation:

```
[local]asr5500(config)# logging disable eventid 96186
```
- To disable NAT attack log generation:

```
[local]asr5500(config)# logging disable eventid 96995
```
- To disable logging for TCP reset message threshold breach:

```
[local]asr5500(config)# logging disable eventid 96203
```
- To disable logging for packets denied by rule:

```
[local]asr5500(config)# logging disable eventid 96159
```

- To disable logging for ICMP unreachable message threshold breach:

```
[local]asr5500(config)# logging disable eventid 96185
```

Increasing Maximum Chunks Per User NAPT for 5G MiFi

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	NAT
Applicable Platform(s)	<ul style="list-style-type: none"> • ASR 5500 • VPC-DI
Feature Default	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>Command Line Interface Reference, Modes C - D Reference Guide</i> • <i>NAT Administration Guide</i>

Revision History

Revision Details	Release
In this release, the NAT inline service is enhanced to support utilization of NAT ports and oversubscription configuration.	21.23

Feature Description

The Network Address Translation (NAT) is enhanced to introduce 5G MiFi device to the 5G network. With the enhancement, when more devices are connected to MiFi, the number of flows increases. The new 5G MiFi device allows up to a large number of devices through WiFi that requires the support of a large number of simultaneous IP flows in the P-GW NAT inline service in each single NAT pool. The NAT inline service is enhanced to support oversubscription and to keep a good utilization of NAT ports.

The **min-port-chunk-per-user** parameter added to the NAT ip-pool configuration guarantees new subscriber to have at least 1 (or n) port-chunks allocation. You can configure NAT pool for each subscriber port-chunk number, when it exceeds the calculated port-chunk number based on either (64k port) or **napt-users-per-ip-address**.

**Note**

- The **min-port-chunk-per-user** is only applicable to NAPT single-ip.
- **min-port-chunk-per-user** and **port-chunk-threshold** are mutually exclusive.
- Allows over subscription configuration.
- The number of port chunks per IP is reduced when you configure port-chunk size to a higher value. This is because the first 1024 ports are reserved.

For more information about Minimum Port-Chunks Reservation and configurations, refer the *Minimum Port Chunks Reservation* section in the *NAT Configuration* chapter of the *NAT Administration Guide*.

Configuring Many-to-One NAT IP Pools

Use the following configuration commands to configure NAT IP pool.

configure

```
context context_name
  ip pool min-port-chunk-per-user max_chunks_per_user
  ip pool port-chunk-threshold
end
```

NOTES:

- **ip pool min-port-chunk-per-user** : Specifies NAT Port minimum number of chunks per user for many-to-one NAT pool. *max_chunks_per_user* must be an integer from 1 through 100.
- **ip pool port-chunk-threshold** and **ip pool min-port-chunk-per-user** are mutually exclusive.

Supported Standards

The NAT feature supports the following RFCs:

- RFC 1631: The IP Network Address Translator (NAT); May 1994
- RFC 1918: Address Allocation for Private Internets; February 1996
- RFC 2663: IP Network Address Translator (NAT) Terminology and Considerations; August 1999
- RFC 2765: Stateless IP/ICMP Translation Algorithm (SIIT); February 2000
- RFC 2766: Network Address Translation - Protocol Translation (NAT-PT); February 2000
- RFC 3022: Traditional IP Network Address Translator (Traditional NAT); January 2001
- RFC 3027: Protocol Complications with the IP Network Address Translator; January 2001
- RFC 3261: SIP: Session Initiation Protocol
- RFC 4787: Network Address Translation (NAT) Behavioral Requirements for Unicast UDP; January 2007
- RFC 4966: Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status; July 2007

- RFC draft-nishitani-cgn-00.txt: Carrier Grade Network Address Translator (NAT) Behavioral Requirements for Unicast UDP, TCP and ICMP; July 2, 2008
- RFC draft-van-beijnum-behave-ftp64-06.txt: IPv6-to-IPv4 translation FTP considerations; May 19, 2009
- RFC draft-ietf-behave-dns64-11.txt: DNS64; February 15, 2010
- RFC draft-ietf-behave-v6v4-xlate-stateful-12.txt: Stateful NAT64; July 10, 2010
- RFC draft-ietf-behave-address-format-10.txt: IPv6 Addressing of IPv4/IPv6 Translators; August 16, 2010
- RFC draft-ietf-behave-v6v4-framework-10.txt: Framework for IPv4/IPv6 Translation; August 17, 2010
- RFC draft-ietf-behave-v6v4-xlate-23.txt: IP/ICMP Translation Algorithm; September 18, 2010
- RFC 6052: IPv6 Addressing of IPv4/IPv6 Translators; October 2010

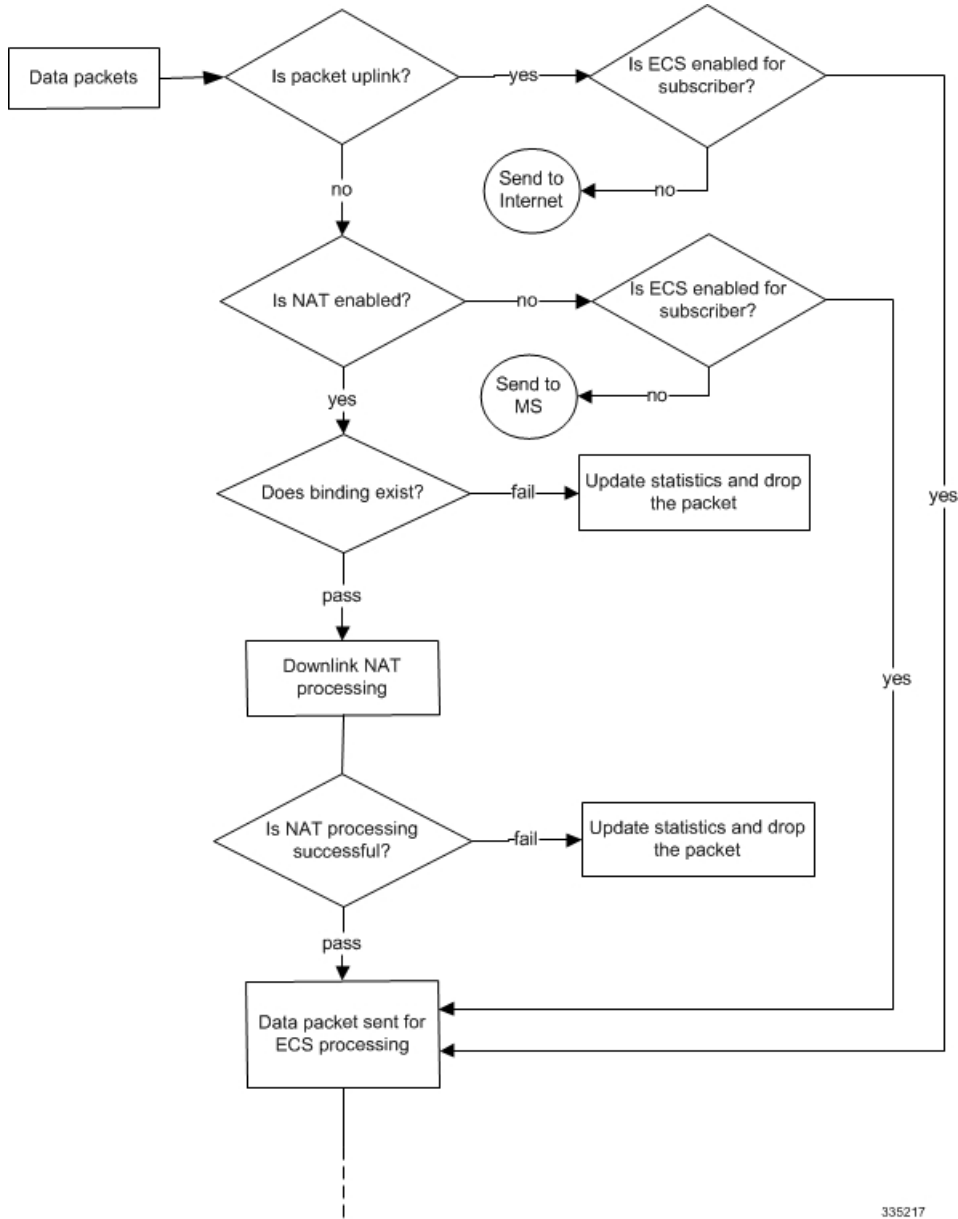
How NAT Works

The following steps describe how NAT works:

-
- Step 1** In the subscriber profile received from the AAA Manager, the SessMgr checks for the following:
- Enhanced Charging Service subsystem must be enabled
 - In the Firewall-and-NAT policy, NAT must be enabled
 - The Firewall-and-NAT policy must be valid
 - For Many-to-One NAT, at least one valid NAT IP pool must be configured in the Firewall-and-NAT policy, and that NAT IP pool must be configured in the context
- Step 2** If all of the above is true, once a private IP address is allocated to the subscriber, the NAT resource to be used for the subscriber is determined. This is only applicable for not-on-demand allocation mode.
- Important** The private IP addresses assigned to subscribers must be from the following ranges for them to get translated: Class A 10.0.0.0 – 10.255.255.255, Class B 172.16.0.0 – 172.31.255.255, Class C 192.168.0.0 – 192.168.255.255, and 100.64.0.0/10 as per RFC 6598
- Important** A subscriber can be allocated only one NAT IP address per NAT IP pool/NAT IP pool group from a maximum of three pools/pool groups. Hence, at any point, there can be a maximum of three NAT IP addresses allocated to a subscriber.
- Step 3** Flow setup is based on the NAT mapping configured for the subscriber:
- In case of one-to-one NAT mapping, the subscriber IP address is mapped to a public IP address. The private source ports do not change. The SessMgr installs a flow using the NAT IP address and a fixed port range (1–65535).
 - In case of many-to-one NAT mapping, a NAT IP address and a port from a port-chunk, are allocated for each connection originating from the subscriber. In order to identify a particular subscriber call line, the SessMgr installs a flow using NAT (public) IP address + NAT ports allocated for the subscriber.


The following figures illustrate the flow of packets in NAT processing.

Figure 3: NAT Processing Flow




335217

Figure 4: ...NAT Processing Flow




335218

Figure 5: ...NAT Processing Flow



335219

Figure 6: ...NAT Processing Flow



335220