



SGSN Pooling

This chapter describes the SGSN Pooling feature.

- [Feature Description, on page 1](#)
- [How it Works, on page 3](#)
- [Configuring the SGSN Pooling feature, on page 10](#)
- [Monitoring and Troubleshooting the SGSN Pooling feature, on page 12](#)

Feature Description

An SGSN pool is a collection of SGSNs configured to serve a common geographical area for a radio network. This common part is referred to as the SGSN pool service. SGSN Pooling is also referred to as Iu/Gb flex support based on if the access is 3G or GPRS respectively.

An SGSN pool provides a flexible and resource-efficient architecture with built-in network redundancy for the GPRS/UMTS packet core network. Each BSC/RNC has the ability to connect to all SGSNs in the pool. If any SGSN becomes unavailable, any terminal attached to that SGSN will be automatically re-routed to another SGSN in the pool by the BSC/RNC. This implies that the SGSN pool provides network level redundancy. SGSN failure is discovered by the BSCs/RNCs and the uplink traffic from the terminal is routed to another SGSN in the pool. The substituting SGSN orders the terminal to re-attach and re-activate any PDP contexts. Therefore service availability is maintained. Please note that all SGSNs in a pool are required to have the same capacity, feature sets and scalability and hence the same vendor, failing which might lead to varying subscriber experience across SGSNs.

In a pooled network, Inter-SGSN routing area updates (RAUs) are avoided and this provides a faster response time, compared to non-pooled networks. With SGSN pool for GPRS/UMTS, Inter-SGSN RAU is replaced by Intra-SGSN RAU, for terminals moving within the pool area. Intra-SGSN RAU provides reduced interruption time for data transfer, compared to Inter-SGSN RAU. Furthermore, due to the fewer Inter-SGSN RAUs, there is less signaling generated on the Gr interface.


When an UE connects to an SGSN in the pool, by Attach or Inter-SGSN RAU (ISRAU) procedures, the UE is allocated a Packet Temporary Mobile Subscriber Identity (P-TMSI) containing a Network Resource Identifier (NRI) identifying the SGSN. The BSC/RNC then identifies the SGSN from the NRI, and routes the user data to the correct SGSN. Load-sharing between the SGSN pool members is thus based on the NRI routing algorithm in the BSC/RNC. UEs that have not yet been assigned a P-TMSI, and MSs without matching NRI, are distributed among the pool members by the BSC/RNC according to the traffic distribution procedure. Once a UE has been allocated a P-TMSI, it stays connected to the same SGSN as long as it remains in the pool area. This period can be quite long, since MSs normally keep the P-TMSI even after power off.

A valid license key is required to enable the SGSN Pooling feature. Contact your Cisco Account or Support representative for information on how to obtain a license.

A Basic Pool Structure

A basic SGSN pool structure is depicted in the diagram below:

Figure 1: A basic pool structure



- Multiple SGSNs form a single logical entity called SGSN pool.
- SGSN pools service areas larger than stand-alone SGSN service areas.
- This set up is compatible with non-pool aware nodes and is transparent to the end-user.

Benefits of SGSN Pooling

1. **Increased Availability:** If one SGSN fails, another SGSN from the pool can substitute it. Any node can be taken out of a pool during maintenance.

2. **Increased Scalability:** More number of SGSN nodes can be added to the pool.
3. **Reduced Signaling:** Number of inter-SGSN routing area updates is reduced.

Pooling Requirements

Listed below are the requirements to support pooling:

1. The SGSN should support configuration of NRI and use that NRI in all the PTMSI issued.
2. If the SGSN is configured as a default SGSN, it should relay SGSN Context Request / Identification request received from peer SGSN (outside of pool) to SGSN (in pool) anchoring that subscriber anchoring SGSN in pool.
3. Support of non-broadcast RAI, null-NRI configurations to allow off-loading of self-SGSN and handle the off-loading of a peer SGSN.

How it Works

P-TMSI - NRI and Coding

Every SGSN is configured with one or several NRIs (O&M). One of these NRIs is part of every Packet temporary Mobile Subscriber identity (P-TMSI) which the SGSN assigns to an UE for connecting via pooled BSC/RNC. For non-pooled BSC/RNC SGSN sets all NRI bits to "0". The P-TMSI allocation mechanism in the SGSN generates P-TMSIs which contain one of the configured NRIs in the relevant bit positions. A NRI has a flexible length up to "6" bits). The maximum number of SGSNs in a pool is limited to "63" (One NRI value reserved for NULL-NRI).

P TMSI is of length "32" bits, where the two top most bits are reserved and always set to "11". The NRI field is included at the beginning of P TMSI starting at bit "23" and down up to bit "18". The most significant bit of the NRI is located at bit "23" of the P TMSI regardless of the configured length of the NRI

Once a subscriber attaches to a new SGSN, a new P-TMSI is allocated by the P-TMSI re-allocation procedure. That P-TMSI contains the NRI of the SGSN. This is also the case when an Inter-SGSN RA update or an Inter-System Change (IRAT) occurs.

Non-Broadcast LAC and RAC

The LAC and RAC information is made available by off-loading the SGSN to the UE in the GMM_ATTCH_ACCEPT/GMM_RAU_ACCEPT message along with the NULL-NRI in the P-TMSI. This value is different from the LAC and RAC that an UE receives from BSS/UTRAN as broadcast information. These parameters are set unique per SGSN node.

SGSN Address Resolution

The following kinds of SGSN address resolution can be identified:

1. Address resolution with NRI.
2. Address resolution without NRI.

Address Resolution with NRI

A NRI based look-up occurs in the following scenarios:

1. An Inter-SGSN RAU occurs within a pooled area. This could be due to one of the SGSNs offloading the subscribers or due to a Gb/ Iu link failure on one of the SGSNs.
2. An Inter-SGSN RAU occurs from a pooled to a non-pooled SGSN. The GTP_SGSN_CONTEXT_REQUEST is routed to the default SGSN in the pool. The default SGSN looks up for the Gn address of the member in the pool based on the NRI retrieved from the P-TMSI in the GTP_SGSN_CONTEXT_REQUEST message received. A local configuration of these entries has to be present in the SGSN Operator Policy.
3. When offloading is enabled, the nb-rai and null-nri of the SGSN which is being offloaded should be configured in the cc-profiles of other SGSN's in the pool. Unless an entry is present, a periodic RAU will not be accepted in the other SGSN's carrying that nb-rai and null-nri. A local configuration of these entries has to be present in the SGSN Operator Policy.

Address resolution without NRI

Address resolution without NRI is used for Inter-SGSN RAU between non-pooled areas or between multiple pools. In this case the SGSN context request is routed towards the default SGSN, which in turn relays the GTP message to the right SGSN based on the NRI value in the P-TMSI.

Refer to the configuration section for the procedure to "Configure an Operator Policy".

Mobility Inside the Pool

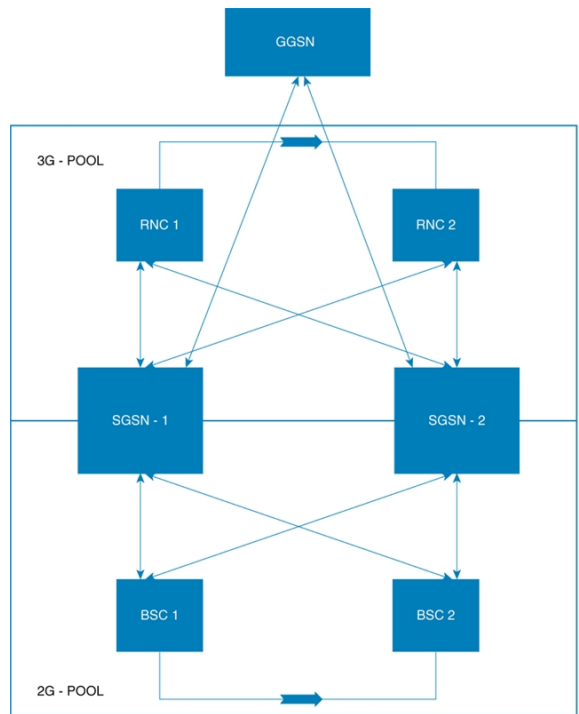
The distribution of UEs in a pool is handled by the BSCs/RNCs.

1. The UE sends an Attach Request or a RA Update Request to a SGSN.
2. This request passes through the BSC/RNC.
3. The BSC/RNC uses the NRI to locate the SGSN.
4. Once the SGSN is located Gb/Iu connection is set up.

If the NRI from the UE is invalid or does not match any of the NRIs of the pool members, the request is directed to one of the pool members by the BSC/RNC. International Mobile Subscriber Identity (IMSI) attaches are also distributed among the SGSN pool members by the BSCs.

Once a P-TMSI containing the NRI of a pool member has been assigned to an UE, the UE stays attached to that pool member as long as it remains in that pool service area. The frequency of inter-SGSN RA updates decreases, as the UE can move over a greater geographical area for one SGSN.

Figure 2: Mobility inside the pool



Consider the scenario depicted in the diagram above:

1. A subscriber attached to SGSN-1 through RNC-1 moves under the coverage area of RNC-2, while being attached to SGSN-1. This results only in an Intra-SGSN RAU.
2. A subscriber attached to SGSN-1 through BSC-1 moves under the coverage area of BSC-2, while being attached to SGSN-1. This results only in an Intra-SGSN RAU.



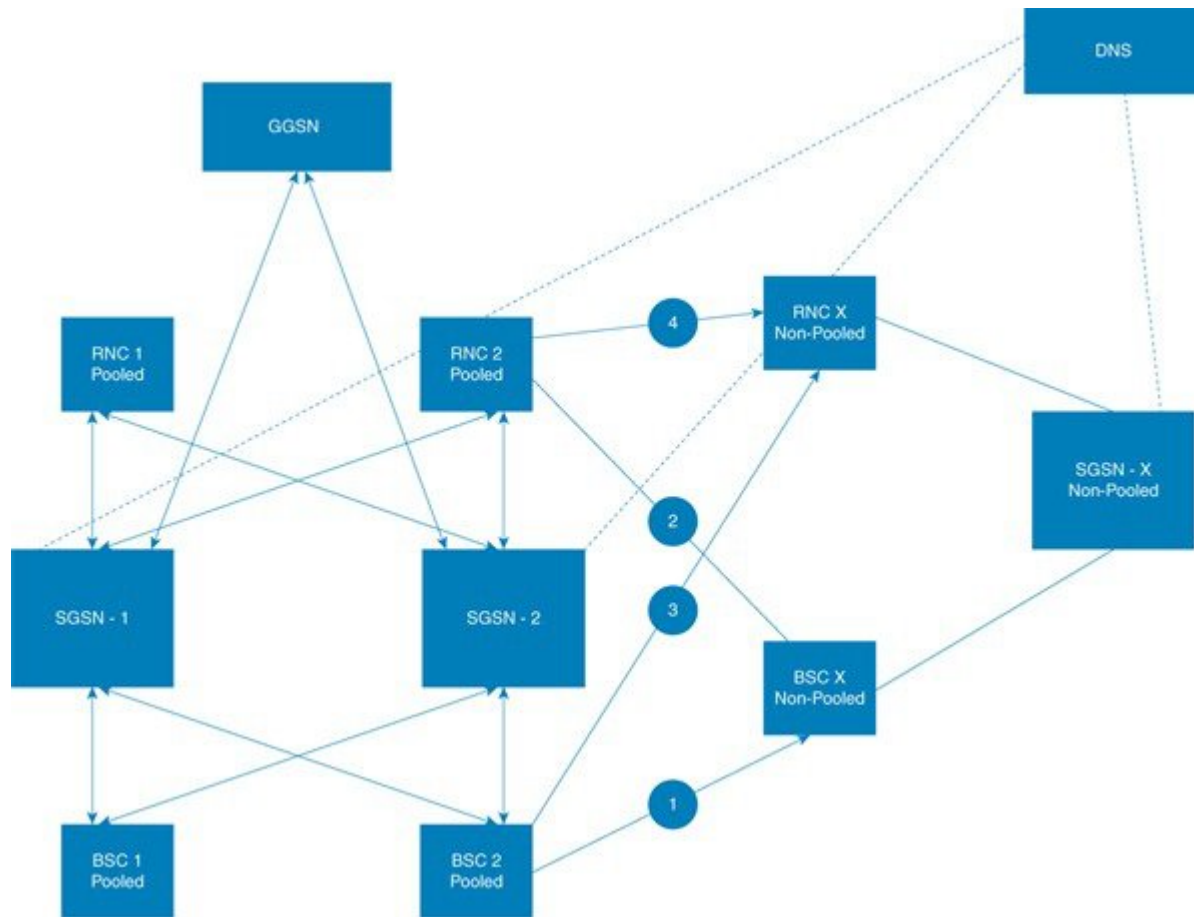
Important

Inter-SGSN RAU within the pool is not very common unless one of the SGSNs within the pool is offloading the subscribers or the Gb/Iu link towards one of the SGSNs from a BSC/RNC is unavailable.

Mobility Outside the Pool

When an UE leaves a pool service area and performs an Attach or a RA update to an SGSN outside the pool service area, the new SGSN cannot identify the old SGSN based on the old Routing Area Identity (RAI). Finding the address of the old SGSN is facilitated by a DNS query with RAI specified. First the new SGSN uses the RAI to identify the default SGSN in the pool. The new SGSN then fetches the subscriber data from the old SGSN and continues with the routing area update procedure.

Figure 3: Mobility outside the Pool



Consider the scenario depicted in the diagram above:

The subscriber movement can be traced through the numbers 1, 2, 3 and 4 in the diagram.

1. The SGSN-X is not pooled. The SGSN-X queries the DNS to identify the source SGSN from where the UE arrived to initiate a GTP_SGSN_CONTEXT_REQUEST.
2. The DNS responds back with the IP address of the default SGSN in the pool, which could be either SGSN-1 or SGSN-2 or both.
3. The address resolution is performed based on the LAC and RAC similar to other Inter-SGSN RAU.
4. The designated default SGSN relays the GTP message to the source SGSN in the pool, which is located using the NRI in the P-TMSI and hence the DNS query with NRI, LAC and RAC.
5. In the implementation above both SGSN-1 and SGSN-2 are designated as default SGSNs to load share the GTP signaling traffic.
6. For every LAC/RAC in the pooled areas the DNS resolves the query into two IP addresses pertaining to the Gn loopback addresses of SGSN-1 and SGSN-2 respectively.

MS Offloading

MS offloading is a procedure of offloading the subscribers from one SGSN in the pool to another SGSN within the same pool. Offloading is performed during the following scenarios:

1. The operator wants to carry out a scheduled maintenance.
2. The operator wants to perform a load re-distribution.
3. To avoid an overload.

Offloading has to be performed with minimum impact on the end users.

Types of MS Offloading:

1. Null-NRI based.
2. Target-NRI based.
3. IMSI based offloading

Null-NRI based Offloading

Null-NRI based offloading is carried out in the following three phases:

Phase - 1

1. UEs performing a RAU or Attach are moved to other SGSN in the pool.
2. When the SGSN receives the Routing Area Update or Attach request, it returns a new P-TMSI with the null-NRI, and non-broadcast LAC and RAC in the accept message.
3. A new Routing Area Update is triggered by setting the periodic routing area update timer to a sufficiently low value in the accept message.
4. The UE sends a new Routing Area Update, the BSC then routes this RAU to a new SGSN due to the presence of a null-NRI. The BSC uses a round robin mechanism to allocate an SGSN for this UE.

Phase - 2

1. All PDP context activation requests are rejected and the UEs are requested to detach and re-attach (Detach request sent from the network with cause code "reattach required").
2. When the UEs re-attach, the SGSN moves them as described above in "Phase 1", that is, by sending the null-NRI and non-broadcast LAC and RAC and triggering a periodic RAU update.

Phase -3

This phase includes scanning through the remaining UEs and initiating a detach procedure for them. The UEs are requested to detach and re-attach, this results in the UEs moving as described in "Phase 1".

UEs being moved from one SGSN can be stopped from registering to the same SGSN again by issuing a CLI command in BSCs connected to the pool. UEs moving into a pool area may also be stopped from registering into a SGSN being off-loaded in the same manner. The move operation will not overload the network, as throttling is supported for both Attach and Inter SGSN RAU procedures.

Target-NRI based offloading

Target NRI based offloading was primarily introduced so that subscribers can be offloaded to a chosen SGSN. In the case of NULL-NRI based offloading there is no control on which SGSN the subscribers are offloaded to. SGSN offloads subscribers by assigning NB-RAI, stamping Target-NRI in PTMSI and reducing periodic routing area update timer during Attach/RAU accept messages.

IMSI-based offloading is carried in the following three phases:

IMSI based offloading

With Target-NRI based method of offloading though there is control on the SGSN to which the subscribers are offloaded, there is no control on the subscribers being offloaded to the SGSN. IMSI-based offloading enhancement allows the operator to choose the subscribers to be offloaded to a particular SGSN.

Phase -1

When a Attach accept or a RAU accept is issued, the offloading configuration is verified and if offloading is enabled, the corresponding NRI is issued (if it is not issued earlier). In case the specific IMSI based offloading configuration is configured, the configured target-nri is used. When offloading is enabled, if ptmsi allocation configuration is absent, a ptmsi is allocated to the subscriber in Attach/RAU accept.

Phase -2

On receiving an activation trigger from the MS, the subscriber is detached and the re-attach required is set to true. The MS will return an attach in due time, after which the MS is offloaded to another SGSN by setting the Target-NRI and NB-RAI appropriately.

Phase -3 0

The subscriber is cleared unconditionally and a detach is sent by setting the re-attach required to true. The subscriber is lost at this stage. In the next attach, the subscriber is offloaded to the configured SGSN.

For information on the procedure to configure MS-Offloading, refer to the section "Configuration of SGSN Pooling - Procedure to configure MS-Offloading".

lu/Gb Flex support over S16/S3 interface

SGSN Pooling support has been extended to S16/S3 interface. The enhancement also includes support for default SGSN functionality for S16/S3 interface as in the case of Gn interface. The peer SGSN in this case is a S4-SGSN. The incoming message (EGTP_CONTEXT_REQ/IDENTIFICATION_REQ) is received from a non-pooled SGSN, it is forwarded to the old-SGSN if the SGSN is configured as default SGSN. The SGSN in a pool is identified on the basis on NRI value and OLD- RAI value. The NRI value is extracted from PTMSI.

Backward compatibility and default SGSN functionality

If a default SGSN that is serving a pool-area receives EGTP signaling it resolves the ambiguity of the multiple SGSNs per RAI by deriving the NRI from the P-TMSI. The SGSN relays the EGTP signaling to the old SGSN identified by the NRI in the old P-TMSI unless the default SGSN itself is the old SGSN. For default-SGSN functionality to work, static IP address entries are mandatory in the call-control profile.

Messages are relayed by the Default-SGSN (Default SGSN functionality and pooling are enabled) in following cases:

- Pooled local RAI and non-local NRI
- Non-local RAI and Null-NRI
- Non-local RAI and Target-NRI

For "Non-local RAI and Null-NRI" and "Non-local RAI and Target-NRI" options, the NB-RAI of other SGSN is considered. It is non-local to the SGSN. No other configuration entries are present at the SGSN other than cc-profile entries.

Mobility Management

The MS performs RA Updates and Attachments, which result in a change of the serving SGSN. In these procedures the new SGSN requests MS specific parameters from the old SGSN. The default SGSN node uses the old RA together with the NRI to derive the signaling address of the old SGSN from its configuration data.

Address and TEID for the Control Plane

- The relaying SGSN forwards the Context Request message to the interface of the old SGSN. The incoming request can arrive over a S3 interface in case of MME or S3 in case of S4-SGSN. However the old RAI interface will be always S16.
- When the default-sgsn relays the message, if the UDP port number is absent in the request received, the default-sgsn adds the "UDP source port number" IE while relaying. This is applicable for both Context Request and Identification Request relay functionality.
- If in an Identification request message, "Address for control plane" is an optional IE. A SGSN within the same SGSN pool with the old SGSN receives the Identification request message it includes the old IP address of the received message in this optional parameter if this IE is not present and relays the message to the old SGSN.
- In cases where default-sgsn has to send a negative response, it sends the message to the IP as indicated in the "S3/S16 Address and TEID for Control Plane" IE and destination port set as indicated by "UDP source port number" IE.
- If an SGSN within the same SGSN pool with the old SGSN receives this message, the SGSN decrements the Hop Counter if this IE is present in the received message. Otherwise, the SGSN includes a Hop Counter with a configured value and relays the message to the old SGSN. This is applicable for both Context Request and Identification Request relay functionality.

For more information refer to 3GPP TS 29.274 (Table 7.3.5-1: Information Elements in a Context Request, Table 7.3.8-1: Information Elements in an Identification Request).

For information on procedure to configure Iu/Gb flex on S16/S3 interface refer to the section "Configuration of SGSN Pooling - Procedure to configure default SGSN (S16/S3 interface)".

Standards Compliance

The SGSN Pooling feature complies with the following standards:

- 3GPP TS 23.236
- 3GPP TS 29.274

Configuring the SGSN Pooling feature

2G-SGSN pool configuration

Listed below are the pre-requisite CLI configurations that should be enabled to configure a 2G SGSN Pool:

1. 2G SGSN Pooling configuration is done under the GPRS service in the Gb context.
2. The NRI value, NRI length, Null-NRI value and non-broadcast LAC/RAC are configured for the GPRS service.
3. The GPRS service is capable of handling both pooled and non-pooled BSCs.

GPRS Service Configuration:

```

config
  context context_name
  gprs-service service_name
  peer-nsei nse_id pooled
  nri length nri_length { nri-value nri_value | null-nri-value null_nri_value non
  broadcast-lac lac_id rac rac_id [ nri-value value ] }
  exit

```

Notes:

- The above configuration must be repeated each time a BSC is added.
- The command **peer-nsei** is used to render a BSC as pooled or non-pooled.

3G-SGSN pool configuration

Listed below are the pre-requisite CLI configurations that should be enabled to configure a 3G SGSN pool:

1. 3G SGSN pooling configuration is done under the IuPS service in the Iu context.
2. The NRI value, NRI length, Null-NRI value and non-broadcast LAC/RAC are configured for the SGSN service.
3. The IuPS service is capable of handling both pooled and non-pooled RNCs.

IuPS Service Configuration

```

config
  context <context_name>
  iups-service <service_name>
  rnc id <rnc_id> pooled
  exit

```

SGSN Service Configuration

```

config
  context <context_name>
  sgsn-service <service_name>

```

```

nri length nri_length [ nri-value nri_value | null-nri-value null_nri_value
non-broadcast mcc mcc mnc mnc lac lac_id rac rac_id nri-value value ]
default nri
no nri
exit

```

To Configure a Default SGSN

This procedure is common to both 2G and 3G SGSN pooling configurations. The SGSN can be configured as a "default SGSN" in the pool under the SGTP service in the Gn context. This configuration is to be performed only once to render a SGSN as a "default SGSN".

```

config
context <context_name>
sgtp-service <service_name>
pool {default-sgsn | hop-counter count}
exit

```

Procedure to Configure a Default SGSN (S16/S3 interface)

The following CLI command under the eGTP Service Configuration mode is used to configure the default SGSN:

```

config
context <context_name>
egtp-service <service_name>
pool {default-sgsn | hop-counter count}
exit

```

The default SGSN receives inbound SGSN context request messages and forwards it to the correct SGSN in the pool based on the NRI bits of the P-TMSI. If the incoming message EGTP_CONTEXT_REQ/IDENTIFICATION_REQ has the hop count IE, the default SGSN decrements the count by one and forwards it to the Old-SGSN. The hop count is not over written even if it is configured. If the hop count IE is missing with incoming message then the then hop count configured gets populated. If no value is configured the default value is chosen. The hop Counter prevents endless relaying of context/identification request. Each relaying SGSN keeps decrementing the hop-counter value if received from the peer-sgsn, otherwise the SGSN includes hop-counter IE. If default-sgsn receives request having hop counter "0", it does not relay the request.

Procedure to Configure an Operator Policy

Step 1:

```

config
operator-policy (default | name policy_name) [-noconfirm]

```

Step 2:

```

config
call-control profile profile_name
sgsn-address { nri nri | rac rac-id lac lac_id | rnc_id rnc_id } [ nri nri
] prefer { fallback-for-dns | local } address { ipv4 ip_address | ipv6
ip_address } interface { gn | s16 }

```

Procedure to Configure MS Offloading

The SGSN **offload** command is used to configure the MS offloading procedure.

The following CLI command (for phase 1 and phase 2 of offloading) is issued for each GPRS/SGSN service:

```
sgsn offload { gprs-service service_name | sgsn-service service_name } {
  activating [ imsi imsi | nri-value nri_value | stop [ imsi imsi | nri-value
    nri_value ] ] | connecting [ nri-value nri_value | stop [ imsi imsi |
  nri-value nri_value | target-nri target_nri ] | t3312-timeout seconds [
  nri-value nri_value | target-nri target_nri ] | target-nri target_nri [ imsi
  imsi | target-count num_to_offload ] }
```



Important

Various combinations of the same command is issued based on whether it is a 2G, 3G, Null-NRI based offloading, Target-NRI based offloading or IMSI based offloading and so on.

The following CLI has to be issued for the phase-3 of offloading:

```
clear subscribers sgsn-service service_name {nri[ <val> | any ]}
```

Consider an SGSN node which was offloaded due to a maintenance requirement, once this SGSN is again operational it will not recover the subscribers attached before the maintenance occurred. In due course this SGSN will be leveraged, with subscribers moved from (partial offload) two or three most loaded SGSNs.

Monitoring and Troubleshooting the SGSN Pooling feature

SGSN Pooling Show Command(s) and/or Outputs

This section provides information regarding show commands and their outputs in support of the SGSN Pooling:

- **show subscribers sgsn-only/gprs-only full all**
- **show sgsn-pool statistics sgsn-service**
- **show sgsn-pool statistics gprs-service**