



Serving Gateway Overview


The Cisco® ASR 5500 core platform provides wireless carriers with a flexible solution that functions as a Serving Gateway (S-GW) in Long Term Evolution-System Architecture Evolution (LTE-SAE) wireless data networks.

- [Product Description, on page 1](#)
- [Network Deployment\(s\), on page 4](#)
- [Features and Functionality - Base Software, on page 9](#)
- [Features and Functionality - Optional Enhanced Feature Software, on page 35](#)
- [How the Serving Gateway Works, on page 44](#)
- [Supported Standards, on page 49](#)

Product Description

The Serving Gateway routes and forwards data packets from the UE and acts as the mobility anchor during inter-eNodeB handovers. Signals controlling the data traffic are received on the S-GW from the MME which determines the S-GW that will best serve the UE for the session. Every UE accessing the EPC is associated with a single S-GW.

Figure 1: S-GW in the Basic E-UTRAN/EPC Network



The S-GW is also involved in mobility by forwarding down link data during a handover from the E-UTRAN to the eHRPD network. An interface from the eAN/ePCF to an MME provides signaling that creates a GRE tunnel between the S-GW and the eHRPD Serving Gateway.

The diagram illustrates the 4G/LTE network architecture. On the left, the E-UTRAN (E-UTRAN) is shown with a mobile phone and an eNodeB. Below it, the eAN (eAN) is shown with a mobile phone and an ePCF. The EPC (E-UTRAN) is a large blue box containing the MME (MME), S-GW (S-GW), P-GW (P-GW), and HSGW (HSGW). The EPCF (eAN) is a large green box containing the eHRPD (eHRPD). The diagram shows the connections between these components and external systems like OCS, PCRF, and 3GPP AAA.

- packet routing and forwarding.
- providing the local mobility anchor (LMA) point for inter-eNodeB handover and assisting the eNodeB reordering function by sending one or more "end marker" packets to the source eNodeB immediately after switching the path.
- mobility anchoring for inter-3GPP mobility (terminating the S4 interface from an SGSN and relaying the traffic between 2G/3G system and a PDN gateway.
- packet buffering for ECM-IDLE mode downlink and initiation of network triggered service request procedure.
- replicating user traffic in the event that Lawful Interception (LI) is required.
- transport level packet marking.
- user accounting and QoS class indicator (QCI) granularity for charging.
- uplink and downlink charging per UE, PDN, and QCI.
- reporting of user location information (ULI).
- support of circuit switched fallback (CSFB) for re-using deployed CS domain access for voice and other CS domain services.

Platform Requirements

The S-GW service runs on a Cisco® ASR 5500 Series chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

Licenses

The S-GW is a licensed Cisco product. Separate session and feature licenses 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*.


Network Deployment(s)

This section describes the supported interfaces and the deployment scenarios of a Serving Gateway.

Serving Gateway in the E-UTRAN/EPC Network


The following figure displays the specific network interfaces supported by the S-GW. Refer to [Supported Logical Network Interfaces \(Reference Points\)](#), on page 5 for detailed information about each interface.

Figure 3: Supported S-GW Interfaces in the E-UTRAN/EPC Network



The following figure displays a sample network deployment of an S-GW, including all of the interface connections with other 3GPP Evolved-UTRAN/Evolved Packet Core network devices.

Figure 4: S-GW in the E-UTRAN/EPC Network



Supported Logical Network Interfaces (Reference Points)

The S-GW provides the following logical network interfaces in support of the E-UTRAN/EPC network:


S1-U Interface

This reference point provides bearer channel tunneling between the eNodeB and the S-GW. It also supports eNodeB path switching during handovers. The S-GW provides the local mobility anchor point for inter-eNodeB hand-overs. It provides inter-eNodeB path switching during hand-overs when the X2 handover interface between base stations cannot be used. The S1-U interface uses GPRS tunneling protocol for user plane (GTP-Uv1). GTP encapsulates all end user IP packets and it relies on UDP/IP transport. The S1-U interface also supports IPSec IKEv2. This interface is defined in 3GPP TS 23.401.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTPv1-U (bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

Figure 5: Supported Protocols on the S1-U Interface




S4 Interface

This reference point provides tunneling and management between the S-GW and a 3GPP S4 SGSN. The interface facilitates soft hand-offs with the EPC network by providing control and mobility support between the inter-3GPP anchor function of the S-GW. This interface is defined in 3GPP TS 23.401.

Supported protocols:

- Transport Layer: UDP
- Tunneling:
 - GTP: IPv4 or IPv6 GTP-C (GTPv2 control/signaling channel) and GTP-U (GTPv1 user/bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

Figure 6: Supported Protocols on the S4 Interface




S5/S8 Interface

This reference point provides tunneling and management between the S-GW and the P-GW, as defined in 3GPP TS 23.401. The S8 interface is an inter-PLMN reference point between the S-GW and the P-GW used during roaming scenarios. The S5 interface is used between an S-GW and P-GW located within the same administrative domain (non-roaming). It is used for S-GW relocation due to UE mobility and if the S-GW needs to connect to a non-collocated P-GW for the required PDN connectivity.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: GTP: GTPv2-C (signaling channel), GTPv1-U (bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

Figure 7: Supported Protocols on the S5/S8 Interface




S11 Interface

This reference point provides GTP-C control signal tunneling between the MME and the S-GW. One GTP-C tunnel is created for each mobile terminal between the MME and S-GW. This interface is defined in 3GPP TS 23.401.

Supported protocols:

- Transport Layer: UDP
- Tunneling: IPv4 or IPv6 GTPv2-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet


Figure 8: Supported Protocols on the S11 Interface

S12 Interface

This reference point provides GTP-U bearer/user direct tunneling between the S-GW and a UTRAN Radio Network Controller (RNC), as defined in 3GPP TS 23.401. This interface provides support for inter-RAT handovers between the 3G RAN and EPC allowing a direct tunnel to be initiated between the RNC and S-GW, thus bypassing the S4 SGSN and reducing latency.

Supported protocols:

- Transport Layer: UDP
- Tunneling: IPv4 or IPv6 GTP-U (GTPv1 bearer/user channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

Figure 9: Supported Protocols on the S12 Interface

Gz Interface


The Gz reference interface enables offline accounting functions on the S-GW. The S-GW collects charging information for each mobile subscriber UE pertaining to the radio network usage. The Gz interface and offline accounting functions are used primarily in roaming scenarios where the foreign P-GW does not support offline charging.

Supported protocols:

- Transport Layer: TCP
- Network Layer: IPv4, IPv6

- Data Link Layer: ARP
- Physical Layer: Ethernet

Figure 10: Supported Protocols on the Gz Interface



Rf Interface



Important


The Rf interface is not supported on the S-GW.

The Diameter Rf interface (3GPP 32.240) is used for offline (post-paid) charging between the Charging Trigger Function (CTF, S-GW) and the Charging Data Function (CDF). It follows the Diameter base protocol state machine for accounting (RFC 3588) and includes support for IMS specific AVPs (3GPP TS 32.299)

Supported protocols:

- Transport Layer: TCP or SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

Figure 11: Supported Protocols on the Rf Interface



Features and Functionality - Base Software

This section describes the features and functions supported by default in the base software for the S-GW service and do not require any additional licenses to implement the functionality.



Important

To configure the basic service and functionality on the system for the S-GW service, refer to the configuration examples provided in the *Serving Gateway Administration Guide*.

3GPP Release 12 Cause-Code IE Support

When an ERAB or a data session is dropped, an operator may need to get, beyond the ULI information, detailed RAN and/or NAS release cause codes information from the access network to be included in the S-GW and P-GW CDRs for call performance analysis, User QoE analysis and proper billing reconciliation. Also, for IMS sessions, the operator may need to get the above information available at P-CSCF.

'Per E-RAB Cause' is received in ERAB Release Command and ER AB Release Indication messages over S1. However RAN and NAS causes are not forwarded to the SGW and PGW, nor provided by the PGW to PCRF.

To resolve this issue a "RAN/NAS Release Cause" information element (IE), which indicates AS and/or NAS causes, has been added to the Session Deletion Request and Delete Bearer Command. The "RAN/NAS Release Cause" provided by the MME is transmitted transparently by the S-GW to the P-GW (if there exists signaling towards P-GW) for further propagation towards the PCRF.

For backward compatibility, the S-GW can still receive the cause code from the CC IE in the S4/S11 messages and/or receive the cause code from some customers' private extension.

Abnormal Bearer Termination Cause in CDR

This feature provides additional information in a S-GW/P-GW CDR for a VoLTE call drop. A dropped bearer was previously reported as a 'abnormalrelease' in the CDR. This feature has the S-GW / P-GW CDRs indicate the proper bearer release for all failure cases identified in the VoLTE Retainability formula. This will provide the customer with the ability to perform gateway/network wide analysis for failures in the network.

New Disconnect reasons are added for GTPC/GTPU path failure and local purge GTPU error indications.

New field abnormalTerminationCause enum 83 is added in the S-GW CDR for a specific customer dictionary.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security. These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the ASR 5500 Platform and an element management system since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

APN-level Traffic Policing

The S-GW now supports traffic policing for roaming scenarios where the foreign P-GW does not enforce traffic classes. Traffic policing is used to enforce bandwidth limitations on subscriber data traffic. It caps packet bursts and data rates at configured burst size and data rate limits respectively for given class of traffic.

Traffic Policing is based on RFC2698- A Two Rate Three Color Marker (trTCM) algorithm. The trTCM meters an IP packet stream and marks its packets green, yellow, or red. A packet is marked red if it exceeds the Peak Information Rate (PIR). Otherwise it is marked either yellow or green depending on whether it exceeds or doesn't exceed the Committed Information Rate (CIR). The trTCM is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.

Backup and Recovery of Key KPI Statistics

Before the Backup and Recovery of Key KPI Statistics feature was implemented, statistics were not backed up and could not be recovered after a SessMgr task restart. Due to this limitation, monitoring the KPI was a problem as the GGSN, P-GW, SAEGW, and S-GW would lose statistical information whenever task restarts occurred.

KPI calculation involves taking a delta between counter values from two time intervals and then determines the percentage of successful processing of a particular procedure in that time interval. When a SessMgr crashes and then recovers, the GGSN, P-GW, SAEGW, and S-GW lose the counter values - they are reset to zero. So, the KPI calculation in the next interval will result in negative values for that interval. This results in a dip in the graphs plotted using the KPI values, making it difficult for operations team to get a consistent view of the network performance to determine if there is a genuine issue or not.

This feature makes it possible to perform reliable KPI calculations even if a SessMgr restart occurs.



Important

For more information on Backup and Recovery of Key KPI Statistics, refer to the *Backup and Recovery of Key KPI Statistics* chapter in this guide.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with an element management system, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a partial list of supported schemas:

- **System:** Provides system-level statistics
- **Card:** Provides card-level statistics
- **Port:** Provides port-level statistics
- **MAG:** Provides MAG service statistics
- **S-GW:** Provides S-GW node-level service statistics

- **IP Pool:** Provides IP pool statistics
- **APN:** Provides Access Point Name statistics

The system supports the configuration of up to four sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

An element management system is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of an element management system parses collected statistics and stores the information in its PostgreSQL database. It can also generate XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, the Bulk Statistics server can archive files to an alternative directory on the server. The directory can be on a local file system or on an NFS-mounted file system on an element management system server.



Important

For more information on bulk statistic configuration, refer to the *Configuring and Maintaining Bulk Statistics* chapter in the *System Administration Guide*.

CDR Support for Including LAPI (Signaling Priority)

This feature is related to M2M support. 3GPP has added the LAPI (signaling priority) indication being sent in the GTP messages, to indicate that the PDN is a low priority bearer and thus can be treated accordingly. APN backoff timer support based on LAPI indication is not yet supported.

3GPP has also added a new AVP in CDR defined in TS 32.298 named "lowPriorityIndicator". If the S-GW receives the LAPI indicator in GTP, the SGW-CDR and generated will contain the LAPI indication.

The benefit of this feature is that it provides support for carrying the LAPI attribute in SGW-CDR and PGW-CDR, so that billing system can then accordingly bill for that PDN.

Circuit Switched Fall Back (CSFB) Support

Circuit Switched Fall Back (CSFB) enables the UE to camp on an EUTRAN cell and originate or terminate voice calls through a forced switchover to the circuit switched (CS) domain or other CS-domain services (for example, Location Services (LCS) or supplementary services). Additionally, SMS delivery via the CS core network is realized without CSFB. Since LTE EPC networks were not meant to directly anchor CS connections, when any CS voice services are initiated, any PS based data activities on the EUTRAN network will be temporarily suspended (either the data transfer is suspended or the packet switched connection is handed over to the 2G/3G network).

CSFB provides an interim solution for enabling telephony and SMS services for LTE operators that do not plan to deploy IMS packet switched services at initial service launch.

The S-GW supports CSFB messaging over the S11 interface over GTP-C. Supported messages are:

- Suspend Notification
- Suspend Acknowledge
- Resume Notification
- Resume Acknowledgement

The S-GW forwards Suspend Notification messages towards the P-GW to suspend downlink data for non-GBR traffic; the P-GW then drops all downlink packets. Later, when the UE finishes with CS services and moves back to E-UTRAN, the MME sends a Resume Notification message to the S-GW which forwards the message to the P-GW. The downlink data traffic then resumes.

Closed Subscriber Group Support

The S-GW supports the following Closed Subscriber Group (CSG) Information Elements (IEs) and Call Detail Record:

- User CSG Information (UCI) IE in S5/S8
- CSG Information Reporting Action IE and functionality in S5/S8
- An SGW-CDR that includes a CSG record

Collision Counter Support in the GTP Layer

GTPv2 message collisions occur in the network when a node is expecting a particular procedure message from a peer node but instead receives a different procedure message from the peer. The S-GW software has been enhanced so that these collisions are now tracked by statistics and handled based on a pre-defined action for each message collision type.

If the SAEGW is configured as a pure P-GW or a pure S-GW, operators will still see the respective collision statistics if they occur.

The output of the **show egtp statistics verbose** command has been enhanced to provide information on GTPv2 message collisions, including:

- **Interface:** The interface on which the collision occurred: SGW (S4/S11), SGW (S5), or PGW (S5).
- **Old Proc (Msg Type):** Indicates the ongoing procedure at eGTP-C when a new message arrived at the interface which caused the collision. The Msg Type in brackets specifies which message triggered this ongoing procedure.
- **New Proc (Msg Type):** The new procedure and message type.
- **Action:** The pre-defined action taken to handle the collision. The action can be one of:
 - No Collision Detected
 - **Suspend Old:** Suspend processing of the original (old) message, process the new message, then resume old message handling.
 - **Abort Old:** Abort the original message handling and processes the new message.
 - **Reject New:** The new message is rejected, and the original (old) message is processed.
 - **Silent Drop New:** Drop the new incoming message, and the old message is processed.
 - **Parallel Hndl:** Both the original (old) and new messages are handled in parallel.
 - **Buffer New:** The new message is buffered and processed once the original (old) message processing is done.
- **Counter:** The number of times each collision type has occurred.

**Important**

The Message Collision Statistics section of the command output only appears if any of the collision statistics have a counter total that is greater than zero.

Sample output:

```
Message Collision Statistics
Interface      Old Proc (Msg Type)      New Proc (Msg Type)      Action      Counter
SGW(S5)        NW Init Bearer Create (95) NW Init PDN Delete (99)  Abort Old      1
```

In this instance, the output states that at the S-GW egress interface (S5) a Bearer creation procedure is going on due to a CREATE BEARER REQUEST(95) message from the P-GW. Before its response comes to the S-GW from the MME, a new procedure PDN Delete is triggered due to a DELETE BEARER REQUEST(99) message from the P-GW.

The action that is carried out due to this collision at eGTP-C is to abort (Abort Old) the Bearer Creation procedure and carry on normally with the PDN Delete procedure. The Counter total of 1 indicates that this collision happened only once.

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establish limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operational thresholds that are configured for the system as described in the *Thresholding Configuration Guide*. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.
- **Port Utilization Thresholds:** If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
- **Port-specific Thresholds:** If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies:** Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

**Important**

For more information on congestion control, refer to the *Congestion Control* chapter in the *System Administration Guide*.

Dedicated Bearer Timeout Support on the S-GW

The S-GW has been enhanced to support a bearer inactivity timeout for GBR and non-GBR S-GW bearer type sessions per Qos Class Identifier (QCI). This enables the deletion of bearers experiencing less data traffic than the configured threshold value. Operators now can configure a bearer inactivity timeout for GBR and non-GBR bearers for more efficient use of system resources.

Downlink Delay Notification

This feature is divided between the following:

- Value Handling
- Throttling
- EPS Bearer ID and ARP Support

Value Handling

This feature provides for the handling of delay value information elements (IEs) at the S-GW. When a delay value is received at the S-GW from a particular MME, the S-GW delays sending data notification requests for all idle calls belonging to that particular MME. Once the timer expires, requests can be sent. The delay value at the S-GW is determined by the factor received in the delay value IE (as a part of either a Modify Bearer Request or a Data Downlink Notification Request) and a hard-coded base factor of 50 ms at the S-GW.

Throttling

This feature provides additional controls allowing the S-GW to set factors that "throttle" the continuous sending and receiving of DDN messages. A single command configures the throttling parameters supporting this feature,

A description of the **ddn throttle** command is located in the S-GW Service Configuration Mode Commands chapter in the *Command Line Interface Reference*.

EPS Bearer ID and ARP Support

This feature allows support for Priority Paging support in the network. This is mainly needed for MPS subscriber support. The paging priority in the paging message is set by MME based on ARP received in Downlink Data Notification message.

In order to support MPS requirement for Priority Paging in the network for MPS subscriber, DDN message has been enhanced to support passing ARP and EBI information. When the S-GW sends a Downlink Data Notification message, it shall include both EPS Bearer ID and ARP. If the Downlink Data Notification is triggered by the arrival of downlink data packets at the S-GW, the S-GW shall include the EPS Bearer ID and ARP associated with the bearer on which the downlink data packet was received. If the Downlink Data Notification is triggered by the arrival of control signaling, the S-GW shall include the EPS Bearer ID and ARP, if present in the control signaling. If the ARP is not present in the control signaling, the S-GW shall include the ARP in the stored EPS bearer context. If multiple EPS Bearers IDs are reported in the Downlink

Data Notification message, the S-GW shall include all the EBI values and the ARP associated with the bearer with the highest priority (lowest ARP value). For more information, see TS 23.401 (section 5.3.4.3) and 29.274 (section 7.2.11). Details are discussed in CR-859 of 3GPP specifications.

DSCP Ingress and Egress and DSCP Marking at the APN Profile

This feature will provide an operator with a configuration to set the DSCP value per APN profile, so different APNs can have different DSCP markings as per QOS requirements for traffic carried by the APN. In addition, the **qci-qos mapping** table is updated with the addition of a **copy-outer** for copying the DSCP value coming in the encapsulation header from the S1u interface to the S5 interface and vice-versa.

Dynamic GTP Echo Timer

The Dynamic GTP Echo Timer enables the eGTP and GTP-U services to better manage GTP paths during network congestion. As opposed to the default echo timer which uses fixed intervals and retransmission timers, the dynamic echo timer adds a calculated round trip timer (RTT) that is generated once a full request/response procedure has completed. A multiplier can be added to the calculation for additional support during congestion periods.

For more information, refer to the *Configuring the GTP Echo Timer* section located in the *Configuring Optional Features on the eGTP S-GW* section of the *Serving Gateway Configuration* chapter.

Event-Based Idle Second Micro-Check Point Generation for the S-GW

Micro-checkpoints were configurable only with the **micro-checkpoint-periodicity** option in the **timeout idle** command in APN Configuration Mode.

The S-GW can be configured to send an idlesec micro-checkpoint from an Active to Standby chassis when the session state changes from active to idle or from idle to active. The micro-checkpoint carries information about the time when the session became active or idle. Upon receipt of the micro-checkpoint, the Standby chassis updates the active/idle time. This process enables the Active and Standby chassis to be synchronized with respect to when a particular session became active or idle.

Since this feature is event-based, it enables the chassis to send micro-checkpoints only when an event occurs, as opposed to sending micro-check points based on a configured time duration, which sends the micro-checkpoints regardless of whether a session state change occurred or not.

To enable this functionality, use the **micro-checkpoint-deemed-idle** keyword in the **timeout idle** command in APN Configuration Mode.

Event Reporting

The S-GW can be configured to send a stream of user event data to an external server. As users attach, detach, and move throughout the network, they trigger signaling events, which are recorded and sent to an external server for processing. Reported data includes failure reasons, nodes selected, user information (IMSI, IMEI, MSISDN), APN, failure codes (if any) and other information on a per PDN-connection level. Event data is used to track the user status via near real time monitoring tools and for historical analysis of major network events.

The *S-GW Event Reporting* chapter at the end of this guide describes the trigger mechanisms and event record elements used for event reporting.

The SGW sends each event record in comma separated values (CSV) format. The record for each event is sent to the external server within 60 seconds of its occurrence. The **session-event-module** command in the Context Configuration mode allows an operator to set the method and destination for transferring event files, as well as the format and handling characteristics of event files. For a detailed description of this command, refer to the *Command Line Interface Reference*.

Idle-mode Signaling Reduction Support

The S-GW now supports Idle-mode Signaling Reduction (ISR) allowing for a control connection to exist between an S-GW and an MME and S4-SGSN. The S-GW stores mobility management parameters from both nodes while the UE stores session management contexts for both the EUTRAN and GERAN/UTRAN. This allows a UE, in idle mode, to move between the two network types without needing to perform racking area update procedures, thus reducing the signaling previously required. ISR support on the S-GW is embedded and no configuration is required however, an optional feature license is required to enable this feature.

ISR support on the S-GW is embedded and no configuration is required, however, an optional feature license must be purchased to enable this feature.

IMSI/IMEI Available in System Event Logs of Type Error and Critical

The S-GW can be configured to provide the IMSI/IMEI in the event log details for the following system event logs of type error and critical, if available. If the IMSI is not available, the S-GW will make a best effort to obtain the IMEI.

Table 1: New and Modified System Event Logs with IMSI/IMEI in System Event Log Details

Event Log #	Description
New Events	
12225	Represents misc_error3 in format "[IMSI <IMSI>] Misc Error3: %s, error code %d"
12226	Represents recover_call_from_crr_failed1 error in format "[IMSI <IMSI>]Sessmgr-%d Recover call from CRR failed for callid:0x%x reason=%s"
12227	Represents aaa_create_session_failed_no_more_sessions1 error in format "[IMSI <IMSI>] Sessmgr-%d Ran out of session handles"
140075	Represents error_log1 in format "[IMSI <IMSI>]%s"
Modified Events	
139001	To print miscellaneous PGW error log.
191006	To print miscellaneous SAEGW error log.
10034	Represents FSM error in format "[IMSI <IMSI>] default call fsm error: ostate=%s(%d) state=%s(%d) event=%s(%d)"
10035	Represents FSM INVALID event in format "[IMSI <IMSI>] default call fsm invalid event: state=%s(%d) event=%s(%d)"

Event Log #	Description
12382	Represents SN_LE_SESSMGR_PGW_REJECT_BEARER_OP in format "[IMSI <IMSI>] Sessmgr-%d: Request to %s bearer rejected. Reason: %s". For example "[IMSI 112233445566778 Sessmgr-1: Request to Create bearer rejected. Reason: Create Bearer Request denied as session recovery is in progress"
12668	Represents fsm_event_error in format "[IMSI <IMSI>] Misc Error: Bad event in sessmgr fsm, event code %d"
12774	Represents pgw_purge_invalid_crr in format "[IMSI <IMSI>] Local %s TEID [%lu] Collision: Clp Connect Time: %lu, Old Clp Callid: %d, Old Clp Connect Time: %lu %s"
12855	Represents ncqos_nrspca_trig_err in format "[IMSI <IMSI>] NCQOS NRSPCA trig rcvd in invalid bcm mode."
12857	Represents ncqos_nrupc_tft_err in format "[IMSI <IMSI>] NCQOS NRUPC Trig : TFT validation failed for nsapi <%u>."
12858	Represents ncqos_nrxx_trig_already in format "[IMSI <IMSI>] NCQOS NRSPCA/NRUPC is already triggered on sess with nsapi <%u>."
12859	Represents ncqos_nrxx_tft_check_fail in format "[IMSI <IMSI>] NCQOS TFT check failed as TFT has invalid opcode for nsapi <%u>:pf_id_bitmap 0x%x and tft_opcode: %d"
12860	Represents ncqos_sec_rej in format "[IMSI <IMSI>] NCQOS Secondary ctxt with nsapi <%u> rejected, due to <%s>."
12861	Represents ncqos_upc_rej in format "[IMSI <IMSI>] UPC Rejected for ctxt with nsapi <%u>, due to <%s>."
12862	Represents ggsn_subsession_invalid_state in format "[IMSI <IMSI>] GGSN subsession invalid state state:<%s>,[event:<%s>]"
11830	Represents gngp_handoff_rejected_for_pdn_ipv4v6 in format "[IMSI <IMSI>] Sessmgr-%d Handoff from PGW-to-GGSN rejected, as GGSN doesnt support Deffered allocation for IPv4v6, dropping the call."
11832	Represents gngp_handoff_rejected_no_non_gbr_bearer_for_def_bearer_selection in format "[IMSI <IMSI>] Sessmgr-%d Handoff from PGW-to-GGSN rejected, as GGSN Callline has no non-GBR bearer to be selected as Default bearer."
11834	Represents gngp_handoff_from_ggsn_rejected_no_ggsn_call in format "[IMSI <IMSI>] Sessmgr-%d Handoff from GGSN-to-PGW rejected, as GGSN call with TEIDC <0x%x> not found."
12960	Represents gtp_pdp_type_mismatch in format "[IMSI <IMSI>] Mismatch between PDP type of APN %s and in create req. Rejecting call"

Event Log #	Description
11282	Represents pcc_intf_error_info in format "[IMSI <IMSI>] %s"
11293	Represents collision_error in format "[IMSI <IMSI>] Collision Error: Temp Failure Handling Delayed Pending Active Transaction: , error code %d"
11917	Represents rcvd_invalid_bearer_binding_req_from_acs in format "[IMSI <IMSI>] Sessmgr %d: Received invalid bearer binding request from ACS."
11978	Represents saegw_uid_error in format "[IMSI <IMSI>] %s"
11994	Represents unwanted_pcc_intf_setup_req error in format "[IMSI <IMSI>] GGSN_INITIATE_SESS_SETUP_REQ is already fwded to PCC interface "
140005	Represents ue_fsm_illegal_event in format "[IMSI <IMSI>] Invalid/unhandled UE event <%s> in state <%s>"
140006	Represents pdn_fsm_illegal_event in format "[IMSI <IMSI>] Invalid/unhandled PDN event <%s> in state <%s>"
140007	Represents epsb_fsm_illegal_event in format "[IMSI <IMSI>] Invalid/unhandled EPSB event <%s> in state <%s>"
10726	Represents saegwdrv_generic_error "[IMSI <IMSI>] %s"

Enable this functionality by using the **logging include-ueid** command in *Global Configuration Mode*. When enabled, the previously mentioned system events of type error and critical will provide the IMSI/IMEI in the logging details, if available.

IP Access Control Lists

IP access control lists allow you to set up rules that control the flow of packets into and out of the system based on a variety of IP packet parameters.

IP access lists, or Access Control Lists (ACLs) as they are commonly referred to, control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria. Once configured, an ACL can be applied to any of the following:

- An individual interface
- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

**Important**

The S-GW supports interface-based ACLs only. For more information on IP access control lists, refer to the *IP Access Control Lists* chapter in the *System Administration Guide*.

IPv6 Capabilities

IPv6 enables increased address efficiency and relieves pressures caused by rapidly approaching IPv4 address exhaustion problem.

The S-GW platform offers the following IPv6 capabilities:

IPv6 Connections to Attached Elements

IPv6 transport and interfaces are supported on all of the following connections:

- Diameter Gxc policy signaling interface
- Diameter Rf offline charging interface
- Lawful Intercept (X1, X2 interfaces)

**Important**

The Diameter Rf offline charging interface is not supported on the S-GW.

Routing and Miscellaneous Features

- OSPFv3
- MP-BGP v6 extensions
- IPv6 flows (Supported on all Diameter QoS and Charging interfaces as well as Inline Services (for example, ECS, P2P detection, Stateful Firewall, etc.)

LIPA Support

A LIPA (Local IP Access) PDN is a PDN Connection for local IP access for a UE connected to a HeNB. The LIPA architecture includes a Local Gateway (LGW) acting as an S-GW GTPv2 peer. The LGW is collocated with HeNB in the operator network behaves as a PGW from SGW perspective. Once the default bearer for the LIPA PDN is established, then data flows directly to the LGW and from there into the local network without traversing the core network of the network operator.

In order to support millions of LIPA GTPC peers, S-GW memory management has been enhanced with regards to GTPv2 procedures and as well as to support the maintenance of statistics per peer node.

Establishment of LIPA PDN follows a normal call flow similar to that of a normal PDN as per 23.401; the specification does not distinguish between a LGW and a PGW call. As a result, the S-GW supports a new configuration option to detect a LIPA peer. As a fallback mechanism, heuristic detection of LIPA peer based on data flow characteristics of a LIPA call is also supported.

Whenever a peer is detected as a LIPA peer, the S-GW will disable GTPC echo mechanism towards that particular peer and stop maintaining some statistics for that peer.

A configuration option in APN profile explicitly indicates that all the PDN's for that APN are LIPA PDN's, so all GTPC peers on S5 for that APN are treated as LGW, and thus no any detection algorithm is applied to detect LGW.

Location Reporting

Location reporting can be used to support a variety of applications including emergency calls, lawful intercept, and charging. This feature reports user location information (ULI).

ULI data reported in GTPv2 messages includes:

- **TAI-ID:** Tracking Area Identity
- **MCC: MNC:** Mobile Country Code, Mobile Network Code
- **TAC:** Tracking Area Code

The S-GW stores the ULI and also reports the information to the accounting framework. This may lead to generation of Gz and Rf Interim records. The S-GW also forwards the received ULI to the P-GW. If the S-GW receives the UE time zone IE from the MME, it forwards this IE towards the P-GW across the S5/S8 interface.

Mapping High Throughput Sessions on Session Managers

Session managers are upgraded to manage several high throughput sessions without sharing the core and without creating a bottleneck on the CPU load.

The gateway – S-GW, SAEGW or P-GW, classifies a session as a high throughput session based on a DCNR flag present in the IE: FLAGS FOR USER PLANE FUNCTION (UPF) SELECTION INDICATION, in the Create Session Request. This DCNR flag is checkpointed and recovered by the gateway.

A high throughput session is placed on a session manager that has no other high throughput session. If all session manager are handling a high throughput session then these sessions are allocated using the Round-Robbin method.



Note

- The selection of session managers for non-high throughput sessions remains the same in the existing setup.
- Non-high throughput sessions are placed along with the high throughput sessions on the same session manager.

Limitations

Managing high throughput sessions on a session manager has the following limitations:

- The following scenarios may result in placing two high throughput sessions on a session manager:
 - Initial attach from eHRPD/2G/3G sessions.
 - IP addresses – both IPv4 and IPv6, are placed on the same session manager.
 - For an S-GW, the second Create Session Request (PDN) from a UE lands directly on a session manager which has the first PDN of the same UE.

- For a collapsed call, the second Create Session Request (PDN) from a UE lands directly on a session manager which has the first PDN of the same UE.
- In a Multi-PDN call from a UE that is capable of DCNR. For example: VoLTE and Internet capable of DCN will be placed on the same session manager.
- The DCNR flag is not defined by 3GPP for Wi-Fi. Therefore, a session cannot be assigned to a session manager during a Wi-Fi to LTE handover with the DCNR flag set.
- This feature manages and supports distribution of high throughput sessions on a session manager but does not guarantee high throughput for a subscriber.
- In some cases, the round robin mechanism could place a high throughput session on a session manager that was already loaded with other high throughput sessions.

MME Restoration Support

MME restoration is a 3GPP specification-based feature designed to gracefully handle the sessions at S-GW once S-GW detects that the MME has failed or restarted. If the S-GW detects an MME failure based on a different restart counter in the Recovery IE in any GTP Signaling message or Echo Request / Response, it will terminate sessions and not maintain any PDN connections.

As a part of this feature, if a S-GW detects that a MME or S4-SGSN has restarted, instead of removing all the resources associated with the peer node, the S-GW shall maintain the PDN connection table data and MM bearer contexts for some specific S5/S8 bearer contexts eligible for network initiated service restoration, and initiate the deletion of the resources associated with all the other S5/S8 bearers.

The S5/S8 bearers eligible for network initiated service restoration are determined by the S-GW based on operator's policy, for example, based on the QCI and/or ARP and/or APN.

The benefit of this feature is that it provides support for the geo-redundant pool feature on the S4-SGSN/MME. In order to restore session when the MME receives a DDN, the S-GW triggers restoration when the serving MME is unavailable, by selecting another MME and sending DDN. This helps in faster service restoration/continuity in case of MME/S4-SGSN failures.

MME Restoration Standards Extension

The solution to recover from MME node failures proposed in the 3GPP standards rely on the deployment of MME pools where each pool services a coverage area. Following an MME failure, the S-GW and MSC/VLR nodes may select the same MME that used to service a UE, if it has restarted, or an alternate MME in the same pool to process Network-initiated signaling that it received in accordance with the NTSR procedures defined in 3GPP TS 23.007 Release 11.

For a failed MME, the S-GW will select an alternate MME from the associated NTSR pool in round robin fashion in each sessmgr instance.

S-GW NTSR Enhancement

When the Network Triggered Service Restoration (NTSR) feature is enabled on the S-GW and it detects an MME failure. If the subscriber served by the failed MME receives any downlink data packets, then the S-GW selects an alternate MME from the NTSR pool in round-robin fashion. The S-GW then sends a Downlink Data Notification (DDN) to the selected MME. This round robin selection of an MME is per sessmgr instance and not system wide.

Previously, operators could configure a maximum of five MME IP addresses in an NTSR pool. To efficiently interoperate with networks containing more than five MMEs, the S-GW has been enhanced so that 10 MME IP addresses can be configured in the NTSR pool. The configured MME IP addresses can be IPv4, IPv6, or a combination of both IPv4 and IPv6.

This feature improves load balancing of DDN messages in the network during an MME failure.

The existing **ntsr-pool** command in *Global Configuration Mode* is used to configure the MME peer IP addresses. The maximum number of MMEs that can be configured has been increased from five to a maximum of 10.

The existing **show ntsr-pool full all** command in Exec Mode is used to view the configured NTSR pool-id, the NTSR Pool type, and the IP addresses of the MME peers. The command output will now show a maximum of 10 MME peer IP addresses.

Multiple PDN Support

Enables an APN-based user experience that enables separate connections to be allocated for different services including IMS, Internet, walled garden services, or offdeck content services.

The Mobile Access Gateway (MAG) function on the S-GW can maintain multiple PDN or APN connections for the same user session. The MAG runs a single node level Proxy Mobile IPv6 (PMIP) tunnel for all user sessions toward the Local Mobility Anchor (LMA) function of the P-GW.

When a user wants to establish multiple PDN connections, the MAG brings up the multiple PDN connections over the same PMIPv6 session to one or more P-GW LMAs. The P-GW in turn allocates separate IP addresses (Home Network Prefixes) for each PDN connection and each one can run one or multiple EPC default and dedicated bearers. To request the various PDN connections, the MAG includes a common MN-ID and separate Home Network Prefixes, APNs and a Handover Indication Value equal to one in the PMIPv6 Binding Updates.



Important

Up to 11 multiple PDN connections are supported.

Node Functionality GTP Echo

This feature helps exchange capabilities of two communicating GTP nodes, and uses the new feature based on whether it is supported by the other node.

This feature allows the S-GW to exchange its capabilities (MABR, PRN, NTSR) with the peer entities through ECHO messages. By this, if both the peer nodes support some common features, then they can make use of new messages to communicate with each other.

With new "node features" IE support in ECHO request/response message, each node can send its supported features (MABR, PRN, NTSR). This way, S-GW can learn the peer node's supported features. S-GW's supported features can be configured by having some configuration at the service level.



Important

Note that the S-GW does not support MABR functionality.

If S-GW wants to use new message, such as P-GW Restart Notification, then S-GW should check if the peer node supports this new feature or not. If the peer does not support it, then S-GW should fall back to old behavior.

If S-GW receives a new message from the peer node, and if S-GW does not support this new message, then S-GW should ignore it. If S-GW supports the particular feature, then it should handle the new message as per the specification.

Online/Offline Charging



Important Offline Charging is not supported on the S-GW.



Important Online Charging is not supported on the S-GW.

The Cisco EPC platforms support offline charging interactions with external OCS and CGF/CDF servers. To provide subscriber level accounting, the Cisco EPC platform supports integrated Charging Transfer Function (CTF) and Charging Data Function (CDF) / Charging Gateway Function (CGF). Each gateway uses Charging-IDs to distinguish between default and dedicated bearers within subscriber sessions.

The ASR 5500 platform offers a local directory to enable temporary file storage and buffer charging records in persistent memory located on a pair of dual redundant RAID hard disks.

The offline charging implementation offers built-in heart beat monitoring of adjacent CGFs. If the Cisco P-GW has not heard from the neighboring CGF within the configurable polling interval, it will automatically buffer the charging records on the local drives until the CGF reactivates itself and is able to begin pulling the cached charging records.

Offline: Gz Reference Interface

The Cisco P-GW and S-GW support 3GPP Release 8 compliant offline charging as defined in TS 32.251, TS 32.297 and 32.298. Whereas the S-GW generates SGW-CDRs to record subscriber level access to PLMN resources, the P-GW creates PGW-CDRs to record user access to external networks. Additionally when Gn/Gp interworking with SGSNs is enabled, the GGSN service on the P-GW records G-CDRs to record user access to external networks.

To provide subscriber level accounting, the Cisco S-GW supports integrated Charging Transfer Function (CTF) and Charging Data Function (CDF). Each gateway uses Charging-IDs to distinguish between default and dedicated bearers within subscriber sessions.

The Gz reference interface between the CDF and CGF is used to transfer charging records via the GTPP protocol. In a standards based implementation, the CGF consolidates the charging records and transfers them via an FTP or SFTP connection over the Bm reference interface to a back-end billing mediation server. The Cisco EPC gateways also offer the ability to transfer charging records between the CDF and CGF serve via FTP or SFTP. CDR records include information such as Record Type, Served IMSI, ChargingID, APN Name, TimeStamp, Call Duration, Served MSISDN, PLMN-ID, etc.

Operator Policy Support

The operator policy provides mechanisms to fine tune the behavior of subsets of subscribers above and beyond the behaviors described in the user profile. It also can be used to control the behavior of visiting subscribers in roaming scenarios, enforcing roaming agreements and providing a measure of local protection against foreign subscribers.

An operator policy associates APNs, APN profiles, an APN remap table, and a call-control profile to ranges of IMSIs. These profiles and tables are created and defined within their own configuration modes to generate sets of rules and instructions that can be reused and assigned to multiple policies. In this manner, an operator policy manages the application of rules governing the services, facilities, and privileges available to subscribers. These policies can override standard behaviors and provide mechanisms for an operator to get around the limitations of other infrastructure elements, such as DNS servers and HSSs.

The operator policy configuration to be applied to a subscriber is selected on the basis of the selection criteria in the subscriber mapping at attach time. A maximum of 1,024 operator policies can be configured. If a UE was associated with a specific operator policy and that policy is deleted, the next time the UE attempts to access the policy, it will attempt to find another policy with which to be associated.

A default operator policy can be configured and applied to all subscribers that do not match any of the per-PLMN or IMSI range policies.

The S-GW uses operator policy to set the Accounting Mode - GTPP (default), RADIUS/Diameter or none. However, the accounting mode configured for the call-control profile will override this setting.

Changes to the operator policy take effect when the subscriber re-attaches and subsequent EPS Bearer activations.

Optimization for egtpinmgr Recovery

Restarting the egtpinmgr task took a significant amount of time for recovery. Hence, the outage time when the GGSN, P-GW, SAEGW, and S-GW were unable to accept any new calls during egtpinmgr recovery was high.

The software is enhanced to optimize the recovery outage window in the event of an egtpinmgr task restart; this has been achieved by optimizing the internal algorithms of egtpinmgr recovery and the data structures required. In addition, recovery time now is dependent only on the number of unique IMSIs and not on the number of sessions for an IMSI.

Peer GTP Node Profile Configuration Support

Provides flexibility to the operators to have different configuration for GTP-C and Lawful Intercept, based on the type of peer or the IP address of the peer

Peer profile feature allows flexible profile based configuration to accommodate growing requirements of customizable parameters with default values and actions for peer nodes of S-GW. With this feature, configuration of GTP-C parameters and disabling/enabling of Lawful Intercept per MCC/MNC or IP address based on rules defined.

A new framework of peer-profile and peer-map is introduced. Peer-profile configuration captures the GTP-C specific configuration and/or Lawful Intercept enable/disable configuration. GTP-C configuration covers GTP-C retransmission (maximum number of retries and retransmission timeout) and GTP echo configuration. Peer-map configuration matches the peer-profile to be applied to a particular criteria. Peer-map supports criteria like MCC/MNC (PLMN-ID) of the peer or IP-address of the peer. Peer-map can then be associated with S-GW service.

Intent of this feature is to provide flexibility to operators to configure a profile which can be applied to a specific set of peers. For example, have a different retransmission timeout for foreign peers as compared to home peers.

P-GW Restart Notification Support

This procedure optimizes the amount of signaling involved on S11/S4 interface when P-GW failure is detected.

P-GW Restart Notification Procedure is a standards-based procedure supported on S-GW to notify detection of P-GW failure to MME/S4-SGSN. P-GW failure detection will be done at S-GW when it detects that the P-GW has restarted (based on restart counter received from the restarted P-GW) or when it detects that P-GW has failed but not restarted (based on path failure detection). When an S-GW detects that a peer P-GW has restarted, it shall locally delete all PDN connection table data and bearer contexts associated with the failed P-GW and notify the MME via P-GW Restart Notification. S-GW will indicate in the echo request/response on S11/S4 interface that the P-GW Restart Notification procedure is supported.

P-GW Restart Notification Procedure is an optional procedure and is invoked only if both the peers, MME/S4-SGSN and S-GW, support it. This procedure optimizes the amount of signaling involved on S11/S4 interface when P-GW failure is detected. In the absence of this procedure, S-GW will initiate the Delete procedure to clean up all the PDNs anchored at that failed P-GW, which can lead to flooding of GTP messages on S11/S4 if there are multiple PDNs using that S-GW and P-GW.

QoS Bearer Management

Provides a foundation for contributing towards improved Quality of User Experience (QoE) by enabling deterministic end-to-end forwarding and scheduling treatments for different services or classes of applications pursuant to their requirements for committed bandwidth resources, jitter and delay. In this way, each application receives the service treatment that users expect.

An EPS bearer is a logical aggregate of one or more Service Data Flows (SDFs), running between a UE and a P-GW in case of GTP-based S5/S8, and between a UE and HSGW in case of PMIP-based S2a connection. An EPS bearer is the level of granularity for bearer level QoS control in the EPC/E-UTRAN. The Cisco P-GW maintains one or more Traffic Flow Templates (TFTs) in the downlink direction for mapping inbound Service Data Flows (SDFs) to EPS bearers. The P-GW maps the traffic based on the downlink TFT to the S5/S8 bearer. The Cisco P-GW offers all of the following bearer-level aggregate constructs:

QoS Class Identifier (QCI): An operator provisioned value that controls bearer level packet forwarding treatments (for example, scheduling weights, admission thresholds, queue management thresholds, link layer protocol configuration, etc). Cisco EPC gateways also support the ability to map the QCI values to DiffServ codepoints in the outer GTP tunnel header of the S5/S8 connection. Additionally, the platform also provides configurable parameters to copy the DSCP marking from the encapsulated payload to the outer GTP tunnel header.



Important

The S-GW does not support non-standard QCI values. QCI values 1 through 9 are standard values and are defined in 3GPP TS 23.203; the S-GW supports these standard values.

Guaranteed Bit Rate (GBR): A GBR bearer is associated with a dedicated EPS bearer and provides a guaranteed minimum transmission rate in order to offer constant bit rate services for applications such as interactive voice that require deterministic low delay service treatment.

Maximum Bit Rate (MBR): The MBR attribute provides a configurable burst rate that limits the bit rate that can be expected to be provided by a GBR bearer (e.g. excess traffic may get discarded by a rate shaping function). The MBR may be greater than or equal to the GBR for a given dedicated EPS bearer.

Aggregate Maximum Bit Rate (AMBR): AMBR denotes a bit rate of traffic for a group of bearers destined for a particular PDN. The Aggregate Maximum Bit Rate is typically assigned to a group of Best Effort service

data flows over the Default EPS bearer. That is, each of those EPS bearers could potentially utilize the entire AMBR, e.g. when the other EPS bearers do not carry any traffic. The AMBR limits the aggregate bit rate that can be expected to be provided by the EPS bearers sharing the AMBR (e.g. excess traffic may get discarded by a rate shaping function). AMBR applies to all Non-GBR bearers belonging to the same PDN connection. GBR bearers are outside the scope of AMBR.

Policing and Shaping: The Cisco S-GW offers a variety of traffic conditioning and bandwidth management capabilities. These tools enable usage controls to be applied on a per-subscriber, per-EPS bearer or per-PDN/APN basis. It is also possible to apply bandwidth controls on a per-APN AMBR capacity. These applications provide the ability to inspect and maintain state for user sessions or Service Data Flows (SDF's) within them using shallow L3/L4 analysis or high touch deep packet inspection at L7. Metering of out-of-profile flows or sessions can result in packet discards or reducing the DSCP marking to Best Effort priority. When traffic shaping is enabled the S-GW enqueues the non-conforming session to the provisioned memory limit for the user session. When the allocated memory is exhausted, the inbound/outbound traffic for the user can be transmitted or policed in accordance with operator provisioned policy.

Removal of Private Extension-based Overcharging Support

Prior to StarOS release 21.0, the Cisco P-GW and S-GW supported the sending and receiving of overcharging protection data via both a non-3GPP Private Extension Information Element (IE), and a 3GPP Indication IE.

However, since 3GPP support to exchange overcharging protection data exists, no operators were using the Overcharging Private Extension (OCP) based solution. It was also reported by some operators that the Private Extension IE carrying overcharging protection data sent by the P-GW was leading to issues at S-GWs of other vendors.

As a result, support for Private Extension-based Overcharging Support is being removed from the Cisco P-GW and S-GW. This has the benefit of preventing unexpected scenarios occurring due to the decoding of a Private Extension ID carrying overcharging protection data at the P-GW/S-GW of other vendors.

Previous and New Behavior for the P-GW

The following table describes the previous and new behavior at the P-GW for Create Session Request (CSReq) and Create Session Response (CSRsp) messages due to the removal of Private Extension Overcharging Support.

Table 2: Previous and New Behavior: CSReq and CSRsp Messages at P-GW Due to Removal of Private Extension Overcharging Support

Scenario No.	IE Carrying OCP Capability Received from S-GW in CSReq	Old Behavior: IE carrying OCP Capability Sent to S-GW in CSRsp	New Behavior: IE Carrying OCP Capability Sent to S-GW in CSRsp
1	Indication IE	Indication IE	No change. Indication IE will be sent in CSRsp.
2	Private Extension IE	Both Private Extension and Indication IEs.	Private Extension IE received from S-GW is ignored. Indication IE is sent in CSRsp.
3	None	Both Private Extension and Indication IEs.	Only Indication IE is sent in CSRsp.

Scenario No.	IE Carrying OCP Capability Received from S-GW in CSReq	Old Behavior: IE carrying OCP Capability Sent to S-GW in CSReq	New Behavior: IE Carrying OCP Capability Sent to S-GW in CSReq
4	Both Private Extension and Indication IEs.	Indication IE	Private Extension IE received from S-GW is ignored. Only Indication IE is sent in CSReq.

The following table describes the previous and new behavior in Modify Bearer Request (MBReq) and Modify Bearer Response (MBRsp) messages due to the removal of Private Extension Overcharging Support.

Table 3: Previous and New Behavior: MBReq and MBRsp Messages at P-GW Due to Removal of Private Extension Overcharging Support

Scenario No.	IE carrying OCP Capability Received from S-GW in MBReq	Old Behavior: IE Carrying OCP Capability Sent to S-GW in MBRsp	New Behavior: IE Carrying OCP Capability Sent to S-GW in MBRsp
1	Indication IE	Indication IE	No Change. Indication IE is sent in MBRsp messages.
2	Private Extension IE	Private Extension IE	Private Extension IE received from S-GW is ignored. Indication IE is sent in MBRsp message.
3	None	Both Private Extension and Indication IEs.	Only the Indication IE is sent in MBRsp message.
4	Both Private Extension and Indication IEs.	Indication IE	Private Extension IE received from the S-GW is ignored. Only the Indication IE is sent in the MBRsp message.

Previous and New Behavior for the S-GW

The following table describes the previous and new behavior in Create Session Response (CSRsp) messages at the S-GW due to the removal of Private Extension Overcharging Support.

Table 4: Previous and New Behavior: CSRsp Messages at the S-GW Due to the Removal of Private Extension Overcharging Support

Scenario No.	Old Behavior: IE Carrying OCP Capability Received from P-GW in CSRsp	New Behavior: IE Carrying OCP Capability Received from PGW in CSRsp	Old Behavior: IE Carrying OCP Capability Sent to MME in CSRsp	New Behavior: IE Carrying OCP Capability Sent to MME in CSRsp
12	Indication IE	No change. OCP capability received as part of the Indication IE is accepted.	Indication IE	No change. Indication IE is sent in CSRsp.

Scenario No.	Old Behavior: IE Carrying OCP Capability Received from P-GW in CSRsp	New Behavior: IE Carrying OCP Capability Received from PGW in CSRsp	Old Behavior: IE Carrying OCP Capability Sent to MME in CSRsp	New Behavior: IE Carrying OCP Capability Sent to MME in CSRsp
2	Private Extension IE	OCP capability received as part of Private Extension IE is ignored.	If gtpc private-extension overcharge-protection is disabled at egtpc service level: Private Extension IE . If gtpc private-extension overcharge-protection is enabled at egtpc service level: Indication IE .	Since the CLI command is deprecated, then the Private Extension IE is forwarded to the MME in CSRsp as would be done for any unknown Private Extension IE.
3	Both Private Extension IE and Indication IE	OCP capability received as part of the Private Extension IE is ignored. Only OCP capability received as a part of the Indication IE is accepted.	If gtpc private-extension overcharge-protection is disabled at egtpc service level: Private Extension IE and Indication IE . If gtpc private-extension overcharge-protection is enabled at egtpc service level: Indication IE .	Since the CLI command is deprecated, then the Private Extension IE is forwarded to the MME in CSRsp as would be done for any unknown Private Extension IE.
4	None	No change.	None	No change.

The following table describes the previous and new behavior in Modify Bearer Response (MBRsp) messages at the S-GW due to the removal of Private Extension Overcharging Support.

Table 5: Previous Behavior and New Behavior: MBRsp Messages at the S-GW Due to the Removal of Private Extension Overcharging Support

Scenario No.	Old Behavior: IE Carrying OCP Capability Received from P-GW in MBRsp	New Behavior: IE Carrying OCP Capability Received from P-GW in MBRsp	Old Behavior: IE Carrying OCP Capability Sent to MME in MBRsp	New Behavior: IE Carrying OCP Capability Sent to MME in MBRsp
1	Indication IE	No change. OCP capability received as part of Indication IE is accepted.	Indication IE	No change. Indication IE is sent in MBRsp.

Scenario No.	Old Behavior: IE Carrying OCP Capability Received from P-GW in MBRsp	New Behavior: IE Carrying OCP Capability Received from P-GW in MBRsp	Old Behavior: IE Carrying OCP Capability Sent to MME in MBRsp	New Behavior: IE Carrying OCP Capability Sent to MME in MBRsp
2	Private Extension IE	OCP capability received as part of Private Extension IE is ignored.	If gtpc private-extension overcharge-protection is disabled at egtpc service level: None . If gtpc private-extension overcharge-protection is enabled at egtpc service level: Indication IE .	Since the CLI command is deprecated, neither one of the two IEs is sent in the MBRsp to the MME for the OCP capability.
3	Both Private Extension ID and Indication IE	OCP capability received as part of the Private Extension IE is ignored. Only the OCP capability received as part of the Indication IE is accepted.	Indication IE	No change. Indication IE is sent in MBRsp.
4	None	None	None	No change.

**Important**

In StarOS releases 21.0 and later, the S-GW will send a MBReq message with only the indication IE for the Pause/Start Charging procedure. The private extension IE is not sent.

**Important**

If the S-GW receives only the private extension IE from the P-GW in the CSRsp/MBRsp message, then the S-GW ignores the private extension IE. As a result, the S-GW assumes that Overcharging Protection is NOT enabled for the P-GW. So, in this scenario, even if the overcharging condition is met at the S-GW, the S-GW will not send a MBReq message for Charging pause to the P-GW.

Rf Diameter Accounting

**Important**

Rf Diameter Accounting is not supported on the S-GW.

Provides the framework for offline charging in a packet switched domain. The gateway support nodes use the Rf interface to convey session related, bearer related or service specific charging records to the CGF and billing domain for enabling charging plans.

The Rf reference interface enables offline accounting functions on the HSGW in accordance with 3GPP Release 8 specifications. In an LTE application the same reference interface is also supported on the S-GW

and P-GW platforms. The Cisco gateways use the Charging Trigger Function (CTF) to transfer offline accounting records via a Diameter interface to an adjunct Charging Data Function (CDF) / Charging Gateway Function (CGF). The HSGW and Serving Gateway collect charging information for each mobile subscriber UE pertaining to the radio network usage while the P-GW collects charging information for each mobile subscriber related to the external data network usage.

The S-GW collects information per-user, per IP CAN bearer or per service. Bearer charging is used to collect charging information related to data volumes sent to and received from the UE and categorized by QoS traffic class. Users can be identified by MSISDN or IMSI.

Flow Data Records (FDRs) are used to correlate application charging data with EPC bearer usage information. The FDRs contain application level charging information like service identifiers, rating groups, IMS charging identifiers that can be used to identify the application. The FDRs also contain the authorized QoS information (QCI) that was assigned to a given flow. This information is used correlate charging records with EPC bearers.

S-GW Collision Handling

GTPv2 message collisions occur in the network when a node is expecting a particular procedure message from a peer node but instead receives a different procedure message from the peer. The S-GW has been enhanced to process collisions at the S-GW ingress interface for:

1. Create Bearer Request or Update Bearer Request messages with Inter-MME/Inter-RAT Modify Bearer Request messages (with and without a ULI change).
2. Downlink Data Notification(DDN) message with Create Bearer Request or Update Bearer Request.

The enhanced behavior is as follows:

1. A CBRReq and MBReq [(Inter MME/Inter RAT (with or without ULI change))] collision at the S-GW ingress interface results in the messages being handled in parallel. The CBRReq will wait for a Create Bearer Response (CBRsp) from the peer. Additionally, an MBReq is sent in parallel to the P-GW.
2. An UBReq and MBReq [(Inter MME/Inter RAT (with or without a ULI change))] collision at the SGW ingress interface is handled with a suspend and resume procedure. The UBReq would be suspended and the MBReq would be processed. Once the MBRsp is sent to the peer from the SGW ingress interface, the UBReq procedure is resumed.
3. The Downlink Data Notification (DDN) message transaction is dis-associated from bearers. So Create Bearer Request (CBR) or Update Bearer Request (UBR) with Downlink Data Notification (DDN) messages are handled parallel.

As a result of this enhancement no S-GW initiated Cause Code message 110 (Temporarily rejected due to handover procedure in progress) will be seen as a part of such collisions. Collisions will be handled in parallel.

Viewing S-GW Collision Statistics

The output of the **show egtpc statistics verbose** command has been enhanced to provide information on GTPv2 message collisions at the S-GW ingress interface, including:

- **Interface:** The interface on which the collision occurred: SGW (S4/S11), SGW (S5).
- **Old Proc (Msg Type):** Indicates the ongoing procedure at eGTP-C when a new message arrived at the interface which caused the collision. The Msg Type in brackets specifies which message triggered this ongoing procedure.
- **New Proc (Msg Type):** The new procedure and message type.
- **Action:** The pre-defined action taken to handle the collision. The action can be one of:
 - **No Collision Detected**

- **Suspend Old:** Suspend processing of the original (old) message, process the new message, then resume old message handling.
- **Abort Old:** Abort the original message handling and processes the new message.
- **Reject New:** The new message is rejected, and the original (old) message is processed.
- **Silent Drop New:** Drop the new incoming message, and the old message is processed.
- **Parallel Hndl:** Both the original (old) and new messages are handled in parallel.
- **Buffer New:** The new message is buffered and processed once the original (old) message processing is done.
- **Counter:** The number of times each collision type has occurred.

**Important**

The *Message Collision Statistics* section of the command output appears only if any of the collision statistics have a counter total that is greater than zero.

S-GW Session Idle Timer

A session idle timer has been implemented on the S-GW to remove stale sessions in those cases where the session is removed on the other nodes but due to some issue remains on the S-GW. Once configured, the session idle timer will tear down those sessions that remain idle for longer than the configured time limit. The implementation of the session idle timer allows the S-GW to more effectively utilize system capacity.

**Important**

The session idle timer feature will not work if the Fast Data Path feature is enabled.

Subscriber Level Trace

Provides a 3GPP standards-based session level trace function for call debugging and testing new functions and access terminals in an LTE environment.

As a complement to Cisco's protocol monitoring function, the S-GW supports 3GPP standards based session level trace capabilities to monitor all call control events on the respective monitored interfaces including S1-U, S11, S5/S8, and Gxc. The trace can be initiated using multiple methods:

- Management initiation via direct CLI configuration
- Management initiation at HSS with trace activation via authentication response messages over S6a reference interface
- Signaling based activation through signaling from subscriber access terminal

Note: Once the trace is provisioned it can be provisioned through the access cloud via various signaling interfaces.

The session level trace function consists of trace activation followed by triggers. The EPC network element buffers the trace activation instructions for the provisioned subscriber in memory using camp-on monitoring. Trace files for active calls are buffered as XML files using non-volatile memory on the local dual redundant hard drives on the ASR 5500 platform. The Trace Depth defines the granularity of data to be traced. Six levels are defined including Maximum, Minimum and Medium with ability to configure additional levels based on vendor extensions.

All call control activity for active and recorded sessions is sent to an off-line Trace Collection Entity (TCE) using a standards-based XML format over an FTP or secure FTP (SFTP) connection. In the current release the IPv4 interfaces are used to provide connectivity to the TCE. Trace activation is based on IMSI or IMEI.

Once a subscriber level trace request is activated it can be propagated via the S5/S8 signaling to provision the corresponding trace for the same subscriber call on the P-GW. The trace configuration will only be propagated if the P-GW is specified in the list of configured Network Element types received by the S-GW. Trace configuration can be specified or transferred in any of the following message types:

- S11: Create Session Request
- S11: Trace Session Activation
- S11: Modify Bearer Request

**Caution**

As subscriber level trace is a CPU intensive activity the maximum number of concurrently monitored trace sessions per Cisco S-GW is 32. Use in a production network should be restricted to minimize the impact on existing services.

Support for One Million S1-U Peers on the S-GW

Due to customer business requirements and production forecasts, support has been added to the StarOS for one million S1-U connections on a single S-GW.

The S1-U interface is the user plane interface carrying user data between an eNodeB and an S-GW received from the terminal. The StarOS now has the capability to scale the number of S1-U peers to one million per VPN context.

A new CLI command has been added to enable operators to set the number of S1-U peers for which statistics should be collected. The limit is restricted to less than one million peers (128k) due to StarOS memory limitations.

How it Works

The gtpumgr uses the following guidelines while allocating peers:

- When a session installation comes from the Session Manager, a peer is created. If statistics are maintained at the Session Manager, the gtpumgr also creates the peer record with the statistics.
- Peer records are maintained per service.
- The number of peers is maintained at the gtpumgr instance level. The limit is one million S1-U peers per gtpumgr instance.
- If the limit of one million peers is exceeded, then peer creation fails. It causes a call installation failure in the gtpumgr, which leads to an audit failure if an audit is triggered.

The feature changes impact all the interfaces/services using the gtpu-service including GGSN/S4-SGSN/SGW/PGW/SAEGW/ePDG/SaMOG/HNB-GW/HeNB-GW for:

- The Gn and Gp interfaces of the General Packet Radio Service (GPRS)
- The Iu, Gn, and Gp interfaces of the UMTS system

- The S1-U, S2a, S2b, S4, S5, S8, and S12 interfaces of the Evolved Packet System (EPS)

Recovery/ICSR Considerations

- After a session manager/gtpmgr recovery or after an ICSR switchover, the same set of peers configured for statistics collection is recovered.
 - Peers with 0 sessions and without statistics are not recovered.
 - Peers with 0 sessions and with statistics are recovered.
 - Peers with Extension Header Support disabled are recovered.
- While upgrading from a previous release, ensure the newer release chassis **gtpu peer statistics threshold** is equal to or greater than the previous release. This ensures that the GTPU peer statistics are preserved during the upgrade. For example, if you are upgrading from release 19.0 to 20.2, and the 19.0 system has 17,000 GTPU sessions, then configure the threshold on the 20.2 chassis to 17,000 as well.

Configuration/Restrictions

- Due to the large number of GTP-U entities connecting to the StarOS, Cisco recommends disabling the GTP-U Path Management feature.
- The configured threshold is not the hard upper limit for statistics allocation because of the distributed nature of system. It is possible that total GTP-U peers with statistics exceeds the configured threshold value to some extent.
- It is assumed that all 1,000,000 peers are not connected to the node in a point-to-point manner. They are connected through routers.
- There will not be any ARP table size change for the StarOS to support this feature.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert:** A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm:** Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps:** SNMP traps have been created that indicate the condition (high threshold crossing and clear) of each of the monitored values.

Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs:** The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated. Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

Logs are supported in both the Alert and the Alarm models.

- **Alarm System:** High threshold alarms generated within the specified polling interval are considered outstanding until a the condition no longer exists or a condition clear alarm is generated. Outstanding alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in an element management system.

The Alarm System is used only in conjunction with the Alarm model.



Important

For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

ULI Enhancements

VoLTE carriers need the last cell/sector updates within the IMS CDRs to assist in troubleshooting customer complaints due to dropped calls as well as LTE network analysis, performance, fraud detection, and operational maintenance. The ultimate objective is to get the last cell sector data in the IMS CDR records in addition to the ULI reporting for session establishment.

To address this issue, the S-GW now supports the following:

- RAN/NAS Cause IE within bearer context of Delete Bearer Command message.
- The S-GW ignores the ULI received as call is going down so there is no point in updating the CDR.

Support for ULI and ULI Timestamp in Delete Bearer Command message had already been added.

Now, when a new ULI is received in the Delete Bearer Command message, a S-GW CDR is initiated.

Features and Functionality - Optional Enhanced Feature Software

This section describes the optional enhanced features and functions for the S-GW service.

Each of the following features require the purchase of an additional license to implement the functionality with the S-GW service.

128k eNodeB Support


128k eNodeB Support is an optional licensed feature. Contact your Cisco account or service representative for detailed information on licensing requirements.

The StarOS supports up to 128k eNodeBs. This is a collective limit for P-GW/S-GW peers in the GTPU data path and S1-U peers. This enables operators to anchor a large number of such eNodeBs or small Cell eNodeBs.

Direct Tunnel

In accordance with standards, one tunnel functionality enables the SGSN to establish a direct tunnel at the user plane level - a GTP-U tunnel, directly between the RAN and the S-GW.

Figure 12: GTP-U with Direct Tunnel



In effect, a direct tunnel reduces data plane latency as the tunnel functionality acts to remove the SGSN from the data plane and limit the SGSN to the control plane for processing. This improves the user experience (for example, expedites web page delivery, reduces round trip delay for conversational services). Additionally, direct tunnel functionality implements the standard SGSN optimization to improve the usage of user plane resources (and hardware) by removing the requirement from the SGSN to handle the user plane processing.

Typically, the SGSN establishes a direct tunnel at PDP context activation using an Update PDP Context Request towards the S-GW. This means a significant increase in control plane load on both the SGSN and S-GW components of the packet core. Hence, deployment requires highly scalable S-GWs since the volume and frequency of Update PDP Context messages to the S-GW will increase substantially. The ASR 5500 platform capabilities ensure control plane capacity will not be a limiting factor with direct tunnel deployment.

**Important**

For more information on direct tunnel support, refer to the *Direct Tunnel for 4G (LTE) Networks* chapter in this guide.

Intelligent Paging for ISR

In case of Idle-mode Signaling Reduction (ISR) active and UE is idle, the S-GW will send Downlink Data Notification (DDN) Message to both the MME and the S4-SGSN if it receives the downlink data or network initiated control message for this UE. In turn, the MME and the S4-SGSN would do paging in parallel consuming radio resources.

To optimize the radio resource, the S-GW will now perform intelligent paging. When configured at S-GW service level for each APN, the S-GW will page in a semi-sequential fashion (one by one to peer MME or S4-SGSN based on last known RAT type) or parallel to both the MME and S4-SGSN.

Inter-Chassis Session Recovery

The ASR 5500 platform provide industry leading carrier class redundancy. The systems protects against all single points of failure (hardware and software) and attempts to recover to an operational state when multiple simultaneous failures occur.

The system provides several levels of system redundancy:

- Under normal N+1 packet processing card hardware redundancy, if a catastrophic packet processing card failure occurs all affected calls are migrated to the standby packet processing card if possible. Calls which cannot be migrated are gracefully terminated with proper call-termination signaling and accounting records are generated with statistics accurate to the last internal checkpoint
- If the Session Recovery feature is enabled, any total packet processing card failure will cause a packet processing card switchover and all established sessions for supported call-types are recovered without any loss of session.

Even though Cisco provides excellent intra-chassis redundancy with these two schemes, certain catastrophic failures which can cause total chassis outages, such as IP routing failures, line-cuts, loss of power, or physical destruction of the chassis, cannot be protected by this scheme. In such cases, the MME Inter-Chassis Session Recovery (ICSR) feature provides geographic redundancy between sites. This has the benefit of not only providing enhanced subscriber experience even during catastrophic outages, but can also protect other systems such as the RAN from subscriber re-activation storms.

ICSR allows for continuous call processing without interrupting subscriber services. This is accomplished through the use of redundant chassis. The chassis are configured as primary and backup with one being active and one in recovery mode. A checkpoint duration timer is used to control when subscriber data is sent from the active chassis to the inactive chassis. If the active chassis handling the call traffic goes out of service, the inactive chassis transitions to the active state and continues processing the call traffic without interrupting the subscriber session. The chassis determines which is active through a propriety TCP-based connection called a redundancy link. This link is used to exchange Hello messages between the primary and backup chassis and must be maintained for proper system operation.

Interchassis Communication

Chassis configured to support ICSR communicate using periodic Hello messages. These messages are sent by each chassis to notify the peer of its current state. The Hello message contains information about the chassis

such as its configuration and priority. A dead interval is used to set a time limit for a Hello message to be received from the chassis' peer. If the standby chassis does not receive a Hello message from the active chassis within the dead interval, the standby chassis transitions to the active state. In situations where the redundancy link goes out of service, a priority scheme is used to determine which chassis processes the session. The following priority scheme is used:

- router identifier
- chassis priority
- chassis MAC address

Checkpoint Messages

Checkpoint messages are sent from the active chassis to the inactive chassis. Checkpoint messages are sent at specific intervals and contain all the information needed to recreate the sessions on the standby chassis, if that chassis were to become active. Once a session exceeds the checkpoint duration, checkpoint data is collected on the session. The checkpoint parameter determines the amount of time a session must be active before it is included in the checkpoint message.



Important

For more information on inter-chassis session recovery support, refer to the *Interchassis Session Recovery* chapter in *System Administration Guide*.

IP Security (IPSec) Encryption

Enables network domain security for all IP packet switched LTE-EPC networks in order to provide confidentiality, integrity, authentication, and anti-replay protection. These capabilities are insured through use of cryptographic techniques.

The Cisco S-GW supports IKEv1 and IPSec encryption using IPv4 addressing. IPSec enables the following two use cases:

- Encryption of S8 sessions and EPS bearers in roaming applications where the P-GW is located in a separate administrative domain from the S-GW
- IPSec ESP security in accordance with 3GPP TS 33.210 is provided for S1 control plane, S1 bearer plane and S1 management plane traffic. Encryption of traffic over the S1 reference interface is desirable in cases where the EPC core operator leases radio capacity from a roaming partner's network.



Important

You must purchase an IPSec license to enable IPSec. For more information on IPSec support, refer to the *IPSec Reference*.

Lawful Intercept

The Cisco Lawful Intercept feature is supported on the S-GW. Lawful Intercept is a licensed-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Layer 2 Traffic Management (VLANs)

Virtual LANs (VLANs) provide greater flexibility in the configuration and use of contexts and services.

VLANs are configured as tags on a per-port basis and allow more complex configurations to be implemented. The VLAN tag allows a single physical port to be bound to multiple logical interfaces that can be configured in different contexts. Therefore, each Ethernet port can be viewed as containing many logical ports when VLAN tags are employed.

**Important**

For more information on VLAN support, refer to the VLANs chapter in the *System Administration Guide*.

New Call Policy for Stale Sessions

Use of new call policy for stale sessions requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

If the newcall policy is set to **reject release-existing-session** and there are pre-existing sessions for the IMSI/IMEI received in Create Session Req, they will be deleted. This allows for no hung sessions on node with newcall policy reject release configured. When S-GW releases the existing call, it follows a proper release process of sending Accounting Stop, sending CCR-T to PCRF/OCS, and generating CDR(s).

New Standard QCI Support

New Standard QCI Support is a license-controlled feature. Contact your Cisco account or support representative for licensing details.

The P-GW/SAEGW/S-GW support additional new 3GPP-defined standard QCIs. QCIs 65, 66, 69, and 70 are now supported for Mission Critical and Push-to-Talk (MC/PTT) applications. These new standard QCIs are supported in addition to the previously supported QCIs of 1 through 9, and operator-defined QCIs 128 through 254.

The StarOS will continue to reject QCIs 10 through 127 sent by the PCRF.

For detailed information on this feature, refer to the *New Standard QCI Support* chapter in this guide.

Overcharging Protection Support

Use of Overcharging Protection requires that a valid license key be installed. Contact your Cisco account representative for information on how to obtain a license.

Overcharging Protection helps in avoiding charging the subscribers for dropped downlink packets while the UE is in idle mode. In some countries, it is a regulatory requirement to avoid such overcharging, so it becomes a mandatory feature for operators in such countries. Overall, this feature helps ensure subscriber are not overcharged while the subscriber is in idle mode.

**Important**

This feature is supported on the P-GW, and S-GW. Overcharging Protection is supported on the SAEGW only if the SAEGW is configured for Pure P or Pure S functionality.

P-GW will never be aware of UE state (idle or connected mode). Charging for downlink data is applicable at P-GW, even when UE is in idle mode. Downlink data for UE may be dropped at S-GW when UE is in idle mode due to buffer overflow or delay in paging. Thus, P-GW will charge the subscriber for the dropped packets, which isn't desired. To address this problem, with Overcharging Protection feature enabled, S-GW will inform P-GW to stop or resume charging based on packets dropped at S-GW and transition of UE from idle to active state.

If the S-GW supports the Overcharging Protection feature, then it will send a CSReq with the PDN Pause Support Indication flag set to 1 in an Indication IE to the P-GW.

If the PGW supports the Overcharging Protection feature then it will send a CSRsp with the PDN Pause Support Indication flag set to 1 in Indication IE and/or private extension IE to the S-GW.

Once the criterion to signal "stop charging" is met, S-GW will send Modify Bearer Request (MBReq) to P-GW. MBReq would be sent for the PDN to specify which packets will be dropped at S-GW. The MBReq will have an indication IE and/or a new private extension IE to send "stop charging" and "start charging" indication to P-GW. For Pause/Start Charging procedure (S-GW sends MBReq), MBRes from P-GW will have indication and/or private extension IE with Overcharging Protection information.

When the MBReq with stop charging is received from a S-GW for a PDN, P-GW will stop charging for downlink packets but will continue sending the packets to S-GW.

P-GW will resume charging downlink packets when either of these conditions is met:

- When the S-GW (which had earlier sent "stop charging" in MBReq) sends "start charging" in MBReq.
- When the S-GW changes (which indicates that maybe UE has relocated to new S-GW).

This feature aligns with the 3GPP TS 29.274: 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C) specification.

For more information on this feature, refer to the *Overcharging Protection Support* chapter in this guide.

Paging Policy Differentiation

This feature requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

S-GW/P-GW provide configuration control to change the DSCP value of the user-datagram packet and outer IP packet (GTP-U tunnel IP header). DSCP marking is done at various levels depending on the configuration. When the Paging Policy Differentiation (PPD) feature is enabled, however, the user-datagram packet DSCP (tunneled IP packet) marking does not change.

Currently, standards specify QCI to DSCP marking of outer GTP-U header only. All configurations present at ECS, P-GW, and S-GW to change the user-datagram packet DSCP value are non-standard. The standards-based PPD feature dictates that P-CSCF or similar Gi entity marks the DSCP of user-datagram packet. This user-datagram packet DSCP value is sent in DDN message by S-GW to MME/S4-SGSN. MME/S4-SGSN uses this DSCP value to give paging priority.



Important

P-GW and S-GW should apply the PPD feature for both Default and Dedicated bearers. As per the specifications, P-GW transparently passes the user-datagram packet towards S-GW. This means, if PPD feature is enabled, operator can't apply different behavior for Default and Dedicated bearers.

**Important**

For more information on paging policy differentiation, refer to the *Paging Policy Differentiation* chapter in this guide.

3GPP Release 12 Load and Overload Support

Use of 3GPP Release 12 (R12) Load and Overload Support requires that a valid license key be installed. Contact your Cisco account representative for information on how to obtain a license.

3GPP R12 GTP-C Load and Overload Control feature is an optional feature which allows a GTP control plane node to send its Load Information to a peer GTP control plane node which the receiving GTP control plane peer node uses to augment existing GW selection procedure for the P-GW and S-GW. Load Information reflects the operating status of the resources of the originating GTP control plane node.

Nodes using GTP control plane signaling may support communication of Overload control Information in order to mitigate overload situation for the overloaded node through actions taken by the peer node(s). This feature is supported over the S5 and S8 interfaces via the GTPv2 control plane protocol.

A GTP-C node is considered to be in overload when it is operating over its nominal capacity resulting in diminished performance (including impacts to handling of incoming and outgoing traffic). Overload control Information reflects an indication of when the originating node has reached such a situation. This information, when transmitted between GTP-C nodes may be used to reduce and/or throttle the amount of GTP-C signaling traffic between these nodes. As such, the Overload control Information provides guidance to the receiving node to decide actions, which leads to mitigation towards the sender of the information.

In brief, load control and overload control can be described in this manner:

- Load control enables a GTP-C entity (for example, an S-GW/P-GW) to send its load information to a GTP-C peer (e.g. an MME/SGSN, ePDG, TWAN) to adaptively balance the session load across entities supporting the same function (for example, an S-GW cluster) according to their effective load. The load information reflects the operating status of the resources of the GTP-C entity.
- Overload control enables a GTP-C entity becoming or being overloaded to gracefully reduce its incoming signaling load by instructing its GTP-C peers to reduce sending traffic according to its available signaling capacity to successfully process the traffic. A GTP-C entity is in overload when it operates over its signaling capacity, which results in diminished performance (including impacts to handling of incoming and outgoing traffic).

A maximum of 64 different load and overload profiles can be configured.

**Important**

3GPP R12 Load and Overload Support is a license-controlled feature. Contact your Cisco representative for more information on licensing requirements.

**Important**

For more information on this feature, refer to the *GTP-C Load and Overload Control Support on the P-GW, SAEGW, and S-GW* chapter in this guide.

3GPP R12 Load and Overload Factor Calculation Enhancement

In capacity testing and also in customer deployments it was observed that the chassis load factor for the 3GPP R12 Load and Overload Support feature was providing incorrect values even when the sessmgr card CPU utilization was high. The root cause is that when the load factor was calculated by taking an average of CPU utilization of sessmgr and demux cards, the demux card CPU utilization never increased more than the sessmgr card CPU utilization. As a result, the system did not go into the overload state even when the sessmgr card CPU utilization was high.

The 3GPP R12 Load/Overload Control Profile feature has been enhanced to calculate the load factor based on the higher value of similar types of cards for CPU load and memory. If the demux card's CPU utilization value is higher than the sessmgr card's CPU utilization value, then the demux card CPU utilization value is used for the load factor calculation.

A new CLI command is introduced to configure different polling intervals for the resource manager so that the demuxmgr can calculate the load factor based on different system requirements.

Operation

The node periodically fetches various parameters (for example, License-Session-Utilization, System-CPU-Utilization and System-Memory-Utilization), which are required for Node level load control information. The node then calculates the load control information itself either based on the weighted factor provided by the user or using the default weighted factor.

Node level load control information is calculated every 30 seconds. The resource manager calculates the system-CPU-utilization and System-Memory-Utilization at a systems level.

For each configured service, load control information can be different. This can be achieved by providing a weightage to the number of active session counts per service license, for example, $((\text{number of active sessions per service} / \text{max session allowed for the service license}) * 100)$.

The node's resource manager calculates the system-CPU-utilization and System-Memory-Utilization at a systems level by averaging CPU and Memory usage for all cards and which might be different from that calculated at the individual card level.

Separate Paging for IMS Service Inspection

Use of Separate Paging for IMS Service Inspection requires that a valid license key be installed. Contact your Cisco account representative for information on how to obtain a license.

When some operators add an additional IMS service besides VoLTE such as RCS, they can use the same IMS bearer between the two services. In this case, separate paging is supported at the MME using an ID which can be assigned from the S-GW according to the services, where the S-GW distinguishes IMS services using a small DPI function to inspect where the traffic comes from using an ID which is assigned from SGW according to the services. The S-GW distinguishes IMS services using a small DPI function to inspect where the traffic comes from (for example IP, Port and so on). After the MME receives this ID from the S-GW after IMS service inspection, the MME will do classified separate paging for each of the services as usual.

Session Recovery Support

Provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

In the telecommunications industry, over 90 percent of all equipment failures are software-related. With robust hardware failover and redundancy protection, any card-level hardware failures on the system can quickly be corrected. However, software failures can occur for numerous reasons, many times without prior indication. StarOS has the ability to support stateful intra-chassis session recovery (ICSR) for S-GW sessions.

When session recovery occurs, the system reconstructs the following subscriber information:

- Data and control state information required to maintain correct call behavior
- Subscriber data statistics that are required to ensure that accounting information is maintained
- A best-effort attempt to recover various timer values such as call duration, absolute time, and others

Session recovery is also useful for in-service software patch upgrade activities. If session recovery is enabled during the software patch upgrade, it helps to preserve existing sessions on the active packet services card during the upgrade process.

**Important**

For more information on session recovery support, refer to the *Session Recovery* chapter in the *System Administration Guide*.

S-GW Paging Enhancements

Use of S-GW Paging requires that a valid license key be installed. Contact your Cisco account representative for information on how to obtain a license.

S-GW Paging includes the following scenarios:

Scenario 1: S-GW sends a DDN message to the MME/S4-SGSN nodes. MME/S4-SGSN responds to the S-GW with a DDN Ack message. While waiting for the DDN Ack message from the MME/S4-SGSN, if the S-GW receives a high priority downlink data, it does not resend a DDN to the MME/S4-SGSN.

Scenario 2: If a DDN is sent to an MME/S4-SGSN and TAU/RAU MBR is received from another MME/S4-SGSN, S-GW does not send DDN.

Scenario 3: DDN is sent to an MME/S4-SGSN and DDN Ack with Cause #110 is received. DDN Ack with cause 110 is treated as DDN failure and standard DDN failure action procedure is initiated.

To handle these scenarios, the following two enhancements have been added to the DDN functionality:

- High Priority DDN at S-GW
- MBR-DDN Collision Handling

These enhancements support the following:

- Higher priority DDN on S-GW and SAEGW, which helps MME/S4-SGSN to prioritize paging.
- Enhanced paging KPI and VoLTE services.
- DDN message and mobility procedure so that DDN is not lost.
- MBR guard timer, which is started when DDN Ack with temporary HO is received. A new CLI command **ddn temp-ho-rejection mbr-guard-timer** has been introduced to enable the guard timer to wait for MBR once the DDN Ack with cause #110 (Temporary Handover In Progress) is received.
- TAU/RAU with control node change triggered DDNs.

In addition to the above functionality, to be compliant with 3GPP standards, support has been enhanced for Downlink Data Notification message and Mobility procedures. As a result, DDN message and downlink data which triggers DDN is not lost. This helps improve paging KPI and VoLTE success rates in scenarios where DDN is initiated because of SIP invite data.

**Important**

For more information on this functionality, refer to the *S-GW Paging Enhancements* chapter in this guide.

How the Serving Gateway Works

This section provides information on the function of the S-GW in an EPC E-UTRAN network and presents call procedure flows for different stages of session setup and disconnect.

The S-GW supports the following network flows:

- [GTP Serving Gateway Call/Session Procedures in an LTE-SAE Network, on page 44](#)

GTP Serving Gateway Call/Session Procedures in an LTE-SAE Network


The following topics and procedure flows are included:

- [Subscriber-initiated Attach \(initial\), on page 44](#)
- [Subscriber-initiated Detach, on page 48](#)

Subscriber-initiated Attach (initial)

This section describes the procedure of an initial attach to the EPC network by a subscriber.

Figure 13: Subscriber-initiated Attach (initial) Call Flow



335262

Table 6: Subscriber-initiated Attach (initial) Call Flow Description

Step	Description
1	The UE initiates the Attach procedure by the transmission of an Attach Request (IMSI or old GUTI, last visited TAI (if available), UE Network Capability, PDN Address Allocation, Protocol Configuration Options, Attach Type) message together with an indication of the Selected Network to the eNodeB. IMSI is included if the UE does not have a valid GUTI available. If the UE has a valid GUTI, it is included.

Step	Description
2	The eNodeB derives the MME from the GUTI and from the indicated Selected Network. If that MME is not associated with the eNodeB, the eNodeB selects an MME using an MME selection function. The eNodeB forwards the Attach Request message to the new MME contained in a S1-MME control message (Initial UE message) together with the Selected Network and an indication of the E-UTRAN Area identity, a globally unique E-UTRAN ID of the cell from where it received the message to the new MME.
3	If the UE is unknown in the MME, the MME sends an Identity Request to the UE to request the IMSI.
4	The UE responds with Identity Response (IMSI).
5	If no UE context for the UE exists anywhere in the network, authentication is mandatory. Otherwise this step is optional. However, at least integrity checking is started and the ME Identity is retrieved from the UE at Initial Attach. The authentication functions, if performed this step, involves AKA authentication and establishment of a NAS level security association with the UE in order to protect further NAS protocol messages.
6	The MME sends an Update Location (MME Identity, IMSI, ME Identity) to the HSS.
7	The HSS acknowledges the Update Location message by sending an Update Location Ack to the MME. This message also contains the Insert Subscriber Data (IMSI, Subscription Data) Request. The Subscription Data contains the list of all APNs that the UE is permitted to access, an indication about which of those APNs is the Default APN, and the EPS subscribed QoS profile for each permitted APN. If the Update Location is rejected by the HSS, the MME rejects the Attach Request from the UE with an appropriate cause.
8	The MME selects an S-GW using the Serving GW selection function and allocates an EPS Bearer Identity for the Default Bearer associated with the UE. If the PDN subscription context contains no P-GW address the MME selects a P-GW as described in clause PDN GW selection function. Then it sends a Create Default Bearer Request (IMSI, MME Context ID, APN, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the selected S-GW.

Step	Description
9	The S-GW creates a new entry in its EPS Bearer table and sends a Create Default Bearer Request (IMSI, APN, S-GW Address for the user plane, S-GW TEID of the user plane, S-GW TEID of the control plane, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the P-GW.
10	If dynamic PCC is deployed, the P-GW interacts with the PCRF to get the default PCC rules for the UE. The IMSI, UE IP address, User Location Information, RAT type, AMBR are provided to the PCRF by the P-GW if received by the previous message.
11	The P-GW returns a Create Default Bearer Response (P-GW Address for the user plane, P-GW TEID of the user plane, P-GW TEID of the control plane, PDN Address Information, EPS Bearer Identity, Protocol Configuration Options) message to the S-GW. PDN Address Information is included if the P-GW allocated a PDN address Based on PDN Address Allocation received in the Create Default Bearer Request. PDN Address Information contains an IPv4 address for IPv4 and/or an IPv6 prefix and an Interface Identifier for IPv6. The P-GW takes into account the UE IP version capability indicated in the PDN Address Allocation and the policies of operator when the P-GW allocates the PDN Address Information. Whether the IP address is negotiated by the UE after completion of the Attach procedure, this is indicated in the Create Default Bearer Response.
12	The Downlink (DL) Data can start flowing towards S-GW. The S-GW buffers the data.
13	The S-GW returns a Create Default Bearer Response (PDN Address Information, S-GW address for User Plane, S-GW TEID for User Plane, S-GW Context ID, EPS Bearer Identity, Protocol Configuration Options) message to the new MME. PDN Address Information is included if it was provided by the P-GW.
14	The new MME sends an Attach Accept (APN, GUTI, PDN Address Information, TAI List, EPS Bearer Identity, Session Management Configuration IE, Protocol Configuration Options) message to the eNodeB.

Step	Description
15	The eNodeB sends Radio Bearer Establishment Request including the EPS Radio Bearer Identity to the UE. The Attach Accept message is also sent along to the UE.
16	The UE sends the Radio Bearer Establishment Response to the eNodeB. In this message, the Attach Complete message (EPS Bearer Identity) is included.
17	The eNodeB forwards the Attach Complete (EPS Bearer Identity) message to the MME.
18	The Attach is complete and UE sends data over the default bearer. At this time the UE can send uplink packets towards the eNodeB which are then tunneled to the S-GW and P-GW.
19	The MME sends an Update Bearer Request (eNodeB address, eNodeB TEID) message to the S-GW.
20	The S-GW acknowledges by sending Update Bearer Response (EPS Bearer Identity) message to the MME.
21	The S-GW sends its buffered downlink packets.
22	After the MME receives Update Bearer Response (EPS Bearer Identity) message, if an EPS bearer was established and the subscription data indicates that the user is allowed to perform handover to non-3GPP accesses, and if the MME selected a P-GW that is different from the P-GW address which was indicated by the HSS in the PDN subscription context, the MME sends an Update Location Request including the APN and P-GW address to the HSS for mobility with non-3GPP accesses.
23	The HSS stores the APN and P-GW address pair and sends an Update Location Response to the MME.
24	Bidirectional data is passed between the UE and PDN.

Subscriber-initiated Detach

This section describes the procedure of detachment from the EPC network by a subscriber.

Figure 14: Subscriber-initiated Detach Call Flow




Table 7: Subscriber-initiated Detach Call Flow Description

Step	Description
1	The UE sends NAS message Detach Request (GUTI, Switch Off) to the MME. Switch Off indicates whether detach is due to a switch off situation or not.
2	The active EPS Bearers in the S-GW regarding this particular UE are deactivated by the MME sending a Delete Bearer Request (TEID) message to the S-GW.
3	The S-GW sends a Delete Bearer Request (TEID) message to the P-GW.
4	The P-GW acknowledges with a Delete Bearer Response (TEID) message.
5	The P-GW may interact with the PCRF to indicate to the PCRF that EPS Bearer is released if PCRF is applied in the network.
6	The S-GW acknowledges with a Delete Bearer Response (TEID) message.
7	If Switch Off indicates that the detach is not due to a switch off situation, the MME sends a Detach Accept message to the UE.
8	The MME releases the S1-MME signaling connection for the UE by sending an S1 Release command to the eNodeB with Cause = Detach.

Supported Standards

The S-GW service complies with some of the standards in the following standards categories:

- [3GPP References, on page 50](#)
- [3GPP2 References, on page 52](#)
- [IETF References, on page 52](#)
- [Object Management Group \(OMG\) Standards, on page 53](#)

3GPP References

Release 12 3GPP References



Important

The S-GW currently supports the following Release 12 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 are listed under *3GPP2 References*.

- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 29.274: 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C); Stage 3
- 3GPP TS 29.281: General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U)

Release 11 3GPP References



Important

The S-GW currently supports the following Release 11 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 are listed under *3GPP2 References*.

- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 29.274: 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C); Stage 3
- 3GPP TS 29.281: General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U)
- 3GPP TS 32.423: Telecommunication management; Subscriber and equipment trace; Trace data definition and management.
- 3GPP TS 36.414: Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 data transport

Release 10 3GPP References



Important

The S-GW currently supports the following Release 10 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 are listed under *3GPP2 References*.

- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 29.274: 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C); Stage 3
- 3GPP TS 29.281: General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U)

Release 9 Supported Standards

- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.060: General Packet Radio Service (GPRS); Service description; Stage 2
- 3GPP TS 23.216: Single Radio Voice Call Continuity (SRVCC); Stage 2 (Release 9)
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 29.274: 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C); Stage 3 (Release 9)
- 3GPP TS 29.281: General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U)
- 3GPP TS 33.106: 3G Security; Lawful Interception Requirements
- 3GPP TS 36.414: Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 data transport

Release 8 Supported Standards

- 3GPP TR 21.905: Vocabulary for 3GPP Specifications
- 3GPP TS 23.003: Numbering, addressing and identification
- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.107: Quality of Service (QoS) concept and architecture
- 3GPP TS 23.203: Policy and charging control architecture
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 23.402: Architecture Enhancements for non-3GPP accesses
- 3GPP TS 23.060: General Packet Radio Service (GPRS); Service description; Stage 2
- 3GPP TS 24.008: Mobile radio interface Layer 3 specification; Core network protocols
- 3GPP TS 24.229: IP Multimedia Call Control Protocol based on SIP and SDP; Stage 3
- 3GPP TS 29.210: Gx application
- 3GPP TS 29.212: Policy and Charging Control over Gx reference point
- 3GPP TS 29.213: Policy and Charging Control signaling flows and QoS
- 3GPP TS 29.214: Policy and Charging Control over Rx reference point

- 3GPP TS 29.274 V8.1.1 (2009-03): 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C); Stage 3 (Release 8)
- 3GPP TS 29.274: Evolved GPRS Tunneling Protocol for Control plane (GTPv2-C), version 8.2.0 (both versions are intentional)
- 3GPP TS 29.275: Proxy Mobile IPv6 (PMIPv6) based Mobility and Tunneling protocols, version 8.1.0
- 3GPP TS 29.281: GPRS Tunneling Protocol User Plane (GTPv1-U)
- 3GPP TS 32.251: Telecommunication management; Charging management; Packet Switched (PS) domain charging
- 3GPP TS 32.295: Charging management; Charging Data Record (CDR) transfer
- 3GPP TS 32.298: Telecommunication management; Charging management; Charging Data Record (CDR) encoding rules description
- 3GPP TS 32.299: Charging management; Diameter charging applications
- 3GPP TS 33.106: 3G Security; Lawful Interception Requirements
- 3GPP TS 36.107: 3G security; Lawful interception architecture and functions
- 3GPP TS 36.300: Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description
- 3GPP TS 36.412: EUTRAN S1 signaling transport
- 3GPP TS 36.413: Evolved Universal Terrestrial Radio Access (E-UTRA); S1 Application Protocol (S1AP)
- 3GPP TS 36.414: Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 data transport

3GPP2 References

- X.P0057-0 v0.11.0 E-UTRAN - eHRPD Connectivity and Interworking: Core Network Aspects

IETF References

- RFC 768: User Datagram Protocol (STD 6).
- RFC 791: Internet Protocol (STD 5).
- RFC 2131: Dynamic Host Configuration Protocol
- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2698: A Two Rate Three Color Marker
- RFC 2784: Generic Routing Encapsulation (GRE)
- RFC 2890: Key and Sequence Number Extensions to GRE
- RFC 3319: Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers

- RFC 3588: Diameter Base Protocol
- RFC 3775: Mobility Support in IPv6
- RFC 3646: DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
- RFC 4006: Diameter Credit-Control Application
- RFC 4282: The Network Access Identifier
- RFC 4283: Mobile Node Identifier Option for Mobile IPv6 (MIPv6)
- RFC 4861: Neighbor Discovery for IP Version 6 (IPv6)
- RFC 4862: IPv6 Stateless Address Autoconfiguration
- RFC 5094: Mobile IPv6 Vendor Specific Option
- RFC 5213: Proxy Mobile IPv6
- Internet-Draft: Proxy Mobile IPv6
- Internet-Draft: GRE Key Option for Proxy Mobile IPv6, work in progress
- Internet-Draft: Binding Revocation for IPv6 Mobility, work in progress

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

