N+2 UP Recovery

Revision History

Revision History

Note Revision history details are not provided for features introduced before release 21.24.

Revision Details	Release
First introduced	Pre 21.24

Feature Description

In accordance with 3GPP, the CP uses Sx-based failure detection which relies on Sx keep alive message responses from the UP.

Using this approach, when the CP does not receive responses from the UP, it retransmits the Sx message a configurable number of times before declaring the UP as down and initiating session tear downs. Depending on the number of retries and the retry interval, the failure detection period can take more than 10 seconds for a reliable determination that the UP is down. Until the Sx-path failure is detected at CP, the CP continues to select the failed-UP and place new PDN-connections from UEs on the failed-UP.

In order to reduce the time it takes for the CP to detect that a UP is down, Cisco CPs can be configured to use the Bidirectional Forwarding Detection (BFD) protocol (RFC 5883 - Bidirectional Forwarding Protocol Detection (BFD) for Multihop Paths).

BFD uses significantly smaller retry periods (in the order of 200 msec) allowing for more rapid UP down detection. It is in addition to the Sx keepalive mechanism for alternate deployment scenarios (e.g. 1:1 UP redundancy).

NOTE: This feature is not dependent on Packet Flow Description (PFD) since PFD pushes common Day-N configurations across the UPs.

Deployment Architecture

This functionality can be enabled only in an "N+2" deployment scenario for UPs that process data sessions. In this scenario, CPs are deployed as an active-standby pair. "N" number of active UPs can be deployed to communicate with the CP. All of these UPs must be part of a specific, non-default, UP group.

NOTE: In N+2, all UPs are active. As such, this functionality only serves to improve data UP recovery times, it is not a redundancy model. It is highly recommended that UPs processing IMS traffic only be deployed in a 1:1 redundancy model.

BFD communications between the CP and UP requires the configuration of one additional loopback IP address per CP/per UP.

Figure 1: BFD Monitoring in N+2 Deployment

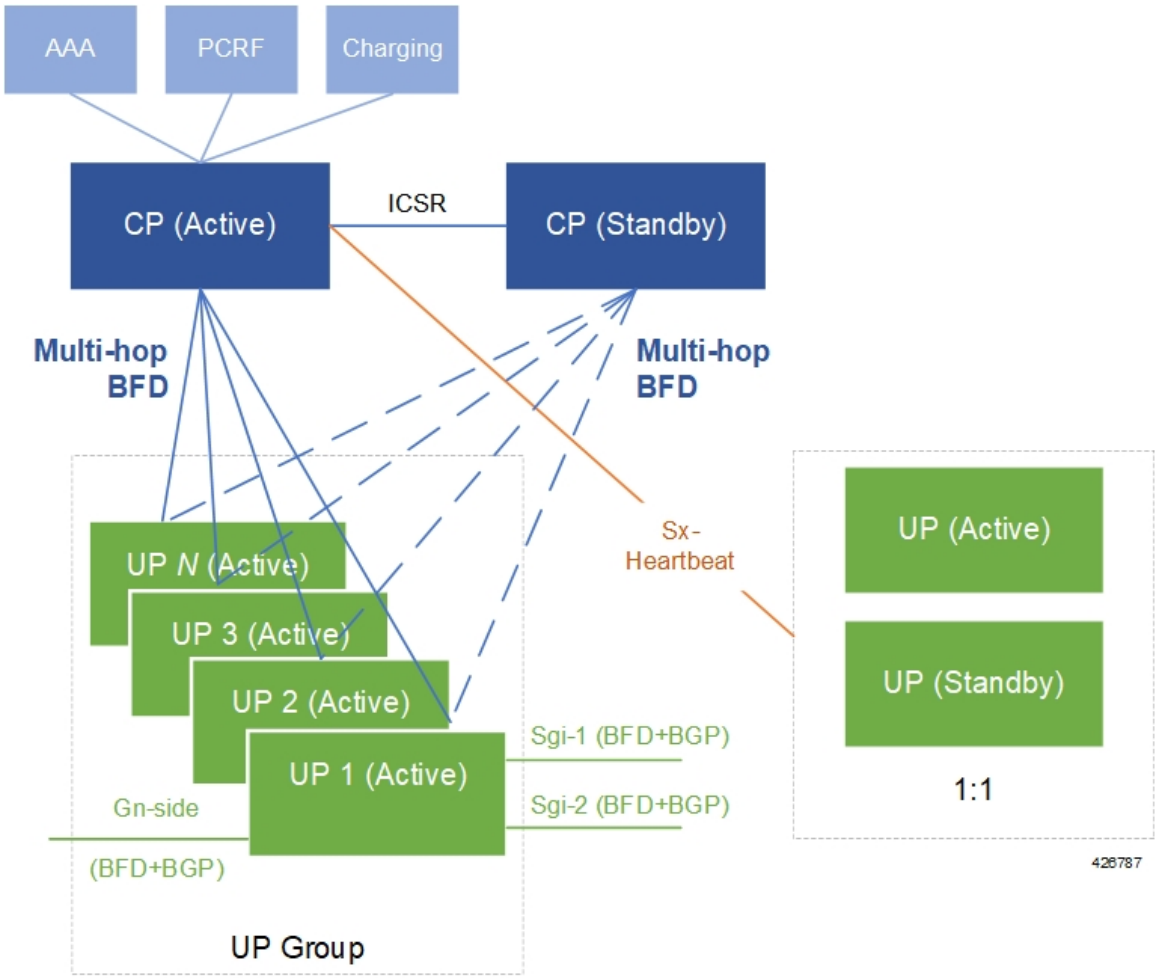

Limitations

- BFD-based CP failure detection is not supported in this release. CP failures can continue to be detected using the existing mechanism of Sx-path failure detection at the UP

  NOTE: It is recommended that Sx-path failure timers be configured more aggressively to more quickly prevent stale UP sessions.

- BGP monitoring on Gi/Gn interface (of UP) is not supported.
- Multi-BFD is not supported.
- BFD must be configured in the same context in which Sx is configured (Gn-side) on both the CP and UP.

How It Works

The figure and the table that follow provide a high-level description of the session detach and re-attach process when a UP is detected as down.

Figure 2: N+2 UP Recovery Flow

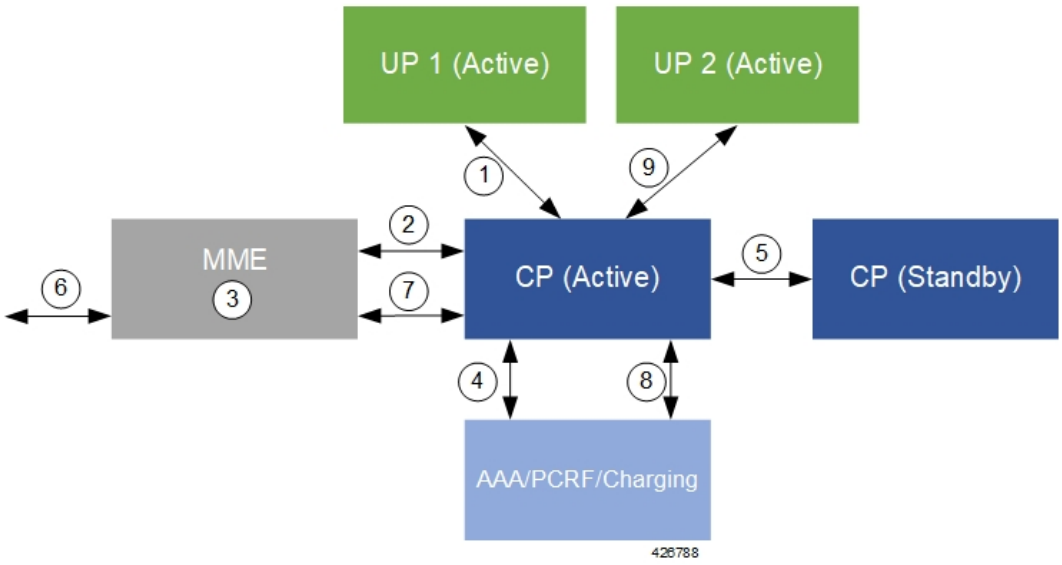

Table 1: N+2 UP Recovery Flow

Number	Description
1	The CP detects a UP failure.
2	The CP sends UP detach session messages to the MME(s) with a cause code of Local-Detach.
3	The MMEs process the request(s) and detach the sessions.
4	The CP communicates with the AAA/PCRF/Charging infrastructure to detach the sessions.
5	The CP (active) communicate with the standby CP to checkpoint the UP detach.

Number	Description
6	UEs whose sessions were previously detached, re-attach to the MME.
7	The MME communicates with the CP to re-attach UE sessions.
8	The CP communicates with the AAA/PCRF/Charging infrastructure to re-attach the sessions.
9	The CP completes the session re-attach process over the Sx interface with an alternate active UP.

Detailed detach and reattach on path failure flows for SAEGW CP/UP, P-GW CP/UP, S-GW CP/UP, and GnGp GGSN CP/UP are in the sections that follow.

Call Flows

SAEGW Detach and Reattach on Path Failure

The figure and the table that follows describe the detach and re-attach on path failure process for SAEGW CPs and UPs.

Figure 3: SAEGW CP/UP Detach and Re-attach on Path Failure Process

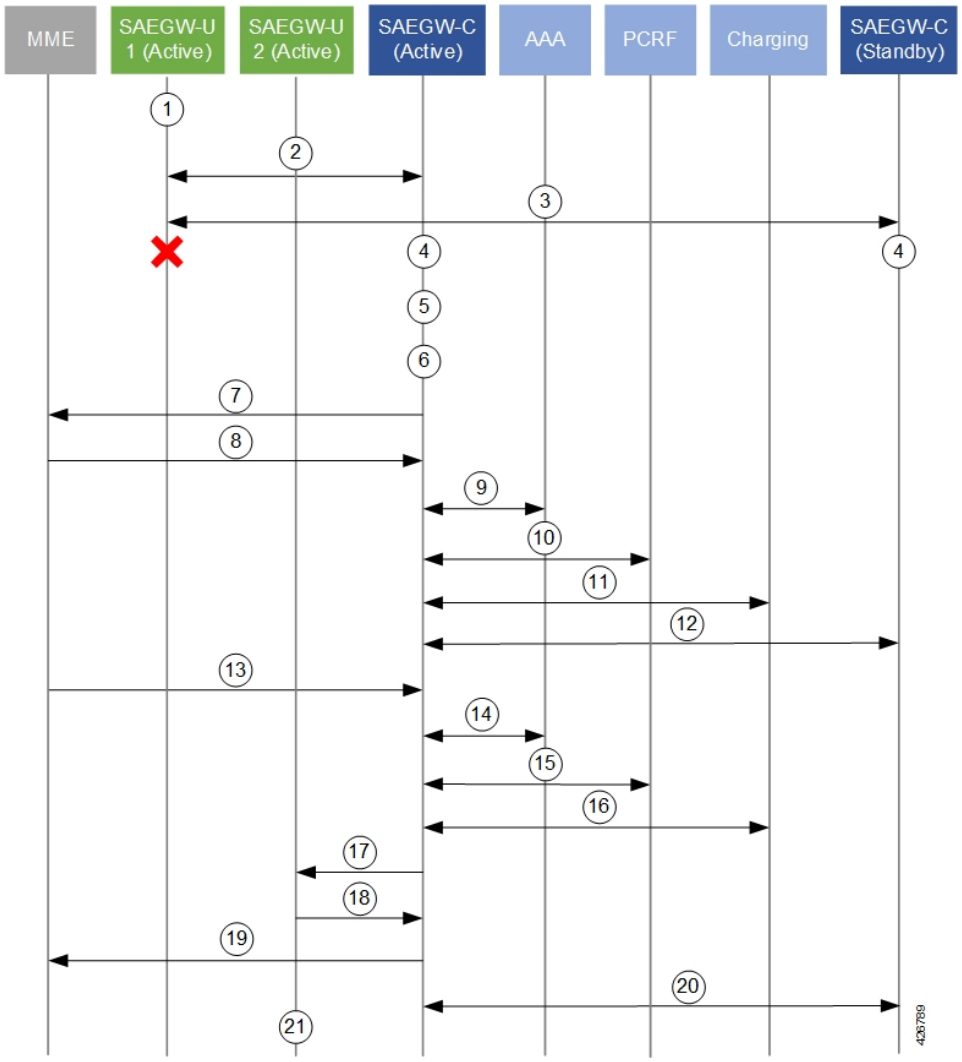

Table 2: SAEGW CP/UP Detach and Re-attach on Path Failure Process

Number	Description
1	UE data sessions are processed by an active SAEGW UP.
2	The active SAEGW CP monitors SAEGW UPs via BFD and Sx-Heartbeat messages.
3	The secondary CP also monitors SAEGW UPs via BFD.
4	The active and standby CPs detect a BFD failure on a UP before eNB detection (relays on Sx timers (interval, retransmission, timeout)).
5	The BFD/VPNMGR on the active CP informs the Sx-demux process of a BFDDown event.
6	The Sx-demux process on the active CP initiates a path Failure notice to all Session Managers on the CP.

Number	Description
7	All Session Managers initiate the process of detaching sessions by sending Delete-bearer-req messages with a cause of Local-Detach to the MME. The detaches are initiated at a pre-defined rate.
8	The MME sends Delete-bearer-resp messages back to the CP. The MME does not page idle UEs with sessions being detached. The MME sends E-RAB release messages to active UEs with sessions being detached.
9	The active CP releases the session release with the AAA server(s).
10	The active CP releases the session with the PCRF.
11	The active CP releases the session with the Charging infrastructure.
12	The active CP syncs session detach information with the secondary CP.
13	For UEs re-initiating their session(s), the MME sends a Create-session-request message to the active CP. The MME selects the CP based on load algorithm (DNS, local config etc.).
14	The active CP processes the session attach request with the AAA server(s).
15	The active CP processes the session attach request with the PCRF.
16	The active CP processes the session attach request with the Charging infrastructure.
17	The active CP sends a Sx Session Establishment Request message to an alternate active UP. The CP selects the UP based on its load algorithm.
18	The UP sends a Sx Session Establishment Response message back to the CP.
19	The CP sends a Create-session-response message to the MME.
20	The active CP syncs information for the newly attached session with the secondary CP.
21	UE data sessions are now processed by the active SAEGW UP.

P-GW Detach and Reattach on Path Failure

The figure and the table that follows describe the detach and re-attach on path failure process for P-GW CPs and UPs.

Figure 4: P-GW CP/UP Detach and Re-attach on Path Failure Process

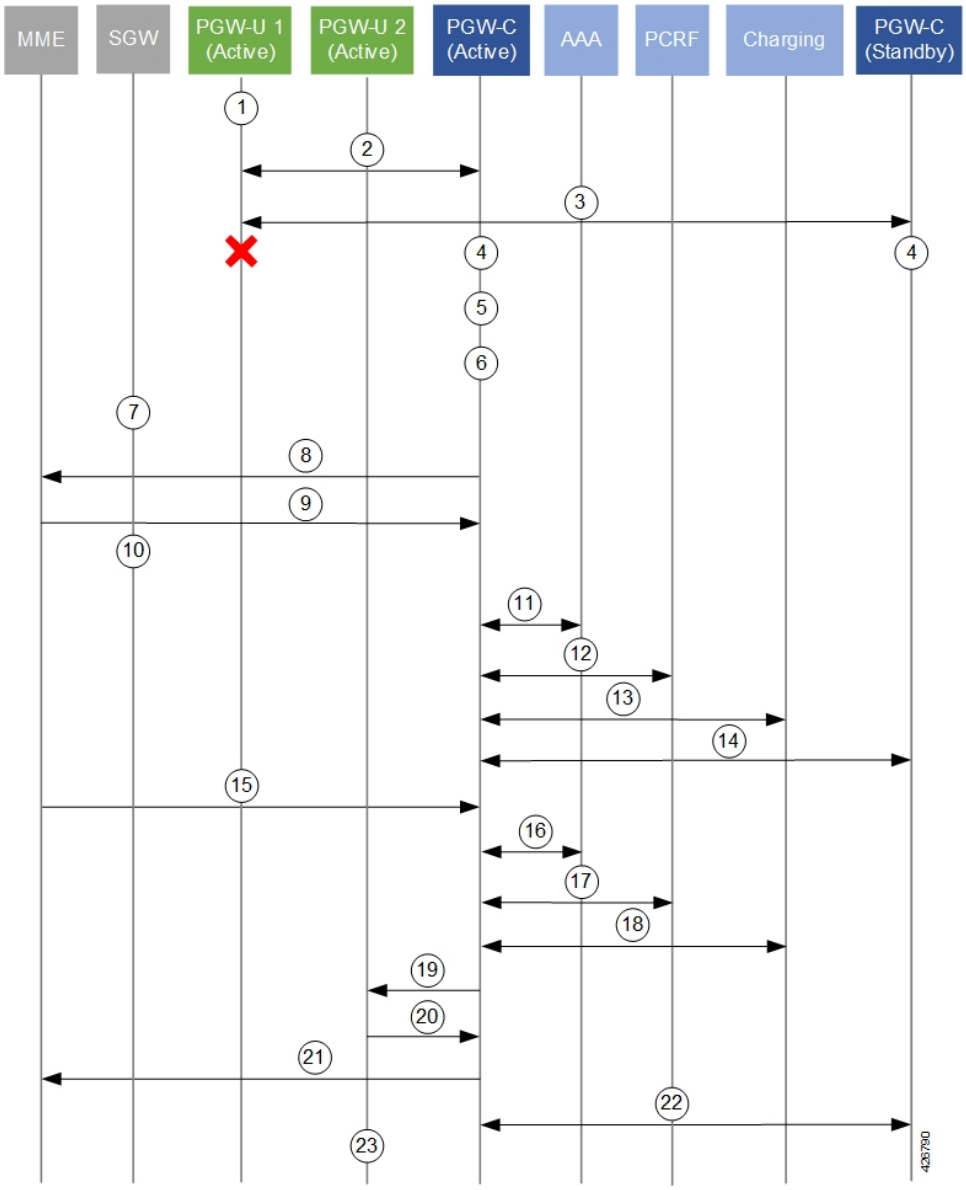

P-GW CP/UP Detach and Re-attach on Path Failure Process

Table 3: P-GW CP/UP Detach and Re-attach on Path Failure Process

Number	Description
1	UE data sessions are processed by an active P-GW UP.
2	The active P-GW CP monitors P-GW UPs via BFD and Sx-Heartbeat messages.
3	The secondary CP also monitors P-GW UPs via BFD.
4	The active and standby CPs detect a BFD failure on a UP before eNB detection (relays on Sx timers (interval, retransmission, timeout)).

Number	Description
5	The BFD/VPNMGR on the active CP informs the Sx-demux process of a BFDDown event.
6	The Sx-demux process on the active CP initiates a path Failure notice to all Session Managers on the CP.
7	The S-GW initiates a db-req to the MME.
8	All Session Managers initiate the process of detach sessions by sending Delete-bearer-req messages with a cause of Local-Detach to the MME. The detaches are initiated at a pre-defined rate.
9	The MME sends Delete-bearer-resp messages back to the CP. The MME does not page idle UEs with sessions being detached. The MME sends E-RAB release messages to active UEs with sessions being detached.
10	The S-GW forwards the db-resp to the PGW-C and removes it's PDN session.
11	The active CP releases the session release with the AAA server(s).
12	The active CP releases the session with the PCRF.
13	The active CP releases the session with the Charging infrastructure.
14	The active CP syncs session detach information with the secondary CP.
15	For UEs re-initiating their session(s), the MME sends a Create-session-request message to the active CP. The MME selects the CP based on load algorithm (DNS, local config etc.).
16	The active CP processes the session attach request with the AAA server(s).
17	The active CP processes the session attach request with the PCRF.
18	The active CP processes the session attach request with the Charging infrastructure.
19	The active CP sends a Sx Session Establishment Request message to an alternate active UP. The CP selects the UP based on its load algorithm.
20	The UP sends a Sx Session Establishment Response message back to the CP.
21	The CP sends a Create-session-response message to the MME.
22	The active CP syncs information for the newly attached session with the secondary CP.
23	UE data sessions are now processed by the active SAEGW UP.

S-GW Detach and Reattach on Path Failure

The figure and the table that follows describe the detach and re-attach on path failure process flow for S-GW CPs and UPs.

Figure 5: S-GW CP/UP Detach and Re-attach on Path Failure Process

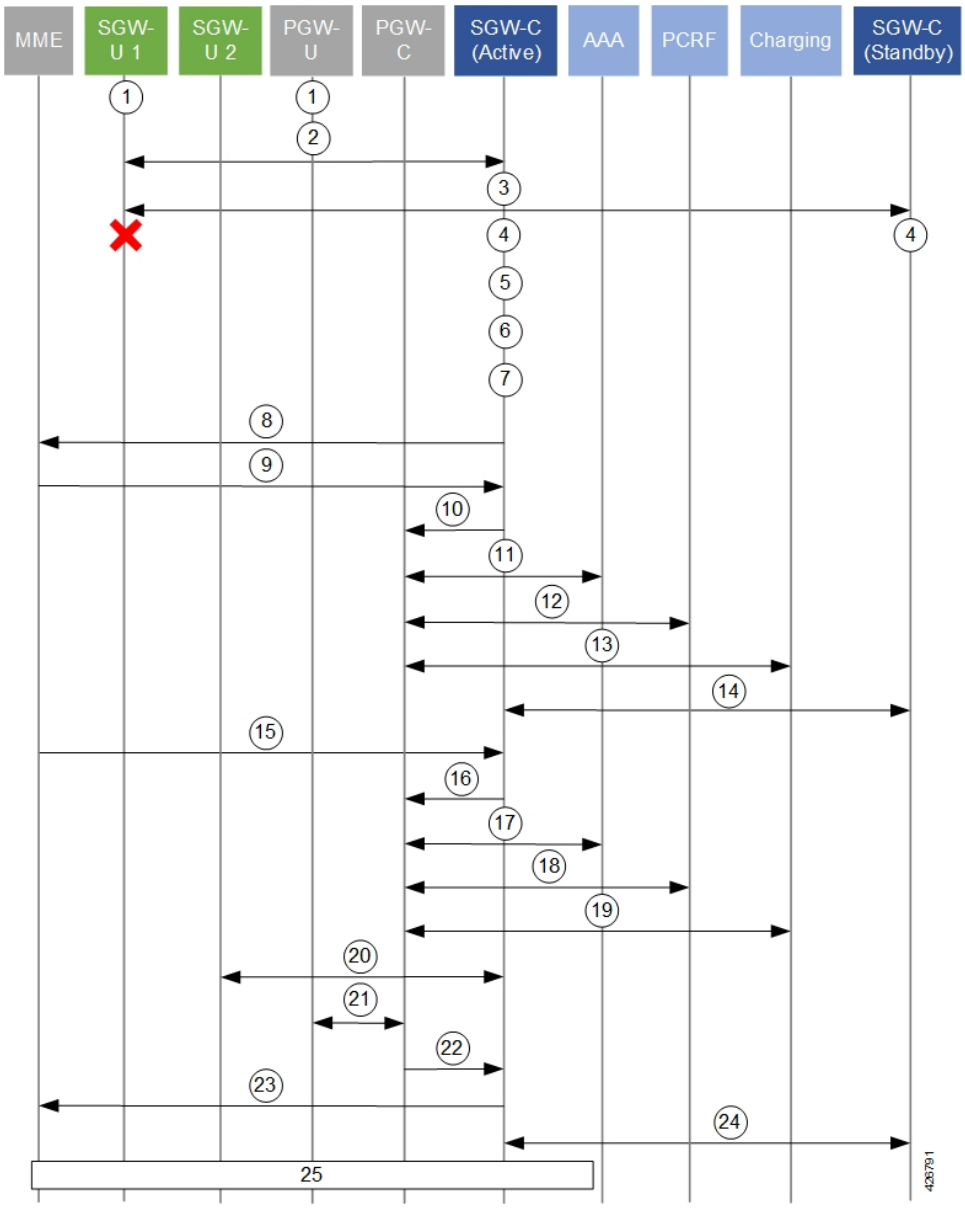

Table 4: S-GW CP/UP Detach and Re-attach on Path Failure Process

Number	Description
1	UE data sessions are processed by an active S-GW UP and an active PGW UP.
2	The active S-GW CP monitors S-GW UPs via BFD and Sx-Heartbeat messages.
3	The secondary S-GW CP also monitors S-GW UPs via BFD.
4	The active and standby S-GW CPs detect a BFD failure on the S-GW UP before eNB detection (relays on Sx timers (interval, retransmission, timeout)).

Number	Description
5	The BFD/VPNMGR on the active S-GW CP informs the Sx-demux process of a BFDDown event.
6	The Sx-demux process on the active CP initiates a path Failure notice to all Session Managers on the CP.
7	The S-GW CP initiates a db-req to the MME.
8	All Session Managers initiate the process of detach sessions by sending Delete-bearer-req messages with a cause of Local-Detach to the MME. The detaches are initiated at a pre-defined rate.
9	The MME sends Delete-bearer-resp messages back to the S-GW CP. The MME does not page idle UEs with sessions being detached. The MME sends E-RAB release messages to active UEs with sessions being detached.
10	The active S-GW CP releases the session release with the PGW UP.
11	The PGW CP releases the session with the AAA server(s).
12	The PGW CP releases the session with the PCRF.
13	The PGW CP releases the session with the Charging infrastructure.
14	The active S-GW CP syncs session detach information with the secondary S-GW CP.
15	For UEs re-initiating their session(s), the MME sends a Create-session-request message to the active S-GW CP. The MME selects the CP based on load algorithm (DNS, local config etc.).
16	The active S-GW CP relays the Create-session-request message to the PGW CP
17	The PGW CP processes the session attach request with the AAA server(s).
18	The PGW CP processes the session attach request with the PCRF.
19	The PGW CP processes the session attach request with the Charging infrastructure.
20	The active S-GW CP exchanges Sx Session Establishment Request and Response messages with an alternate active S-GW UP.
21	The active PGW CP exchanges Sx Session Establishment Request and Response messages with an active PGW UP.
22	The PGW CP sends a Create-session-response message to the S-GW CP.
23	The S-GW CP sends a Create-session-response message to the MME.
24	The active S-GW CP syncs information for the newly attached session with the secondary S-GW CP.
25	The S-GW CP and the complete the Modify Bearer Request procedure with the MME before UE data can flow through the active UPs.

GnGp GGSN Detach and Reattach on Path Failure

The figure and the table that follows describe the detach and re-attach on path failure process flow for GnGp GGSN CPs and UPs.

Figure 6: GnGp GGSN CP/UP Detach and Re-attach on Path Failure Process

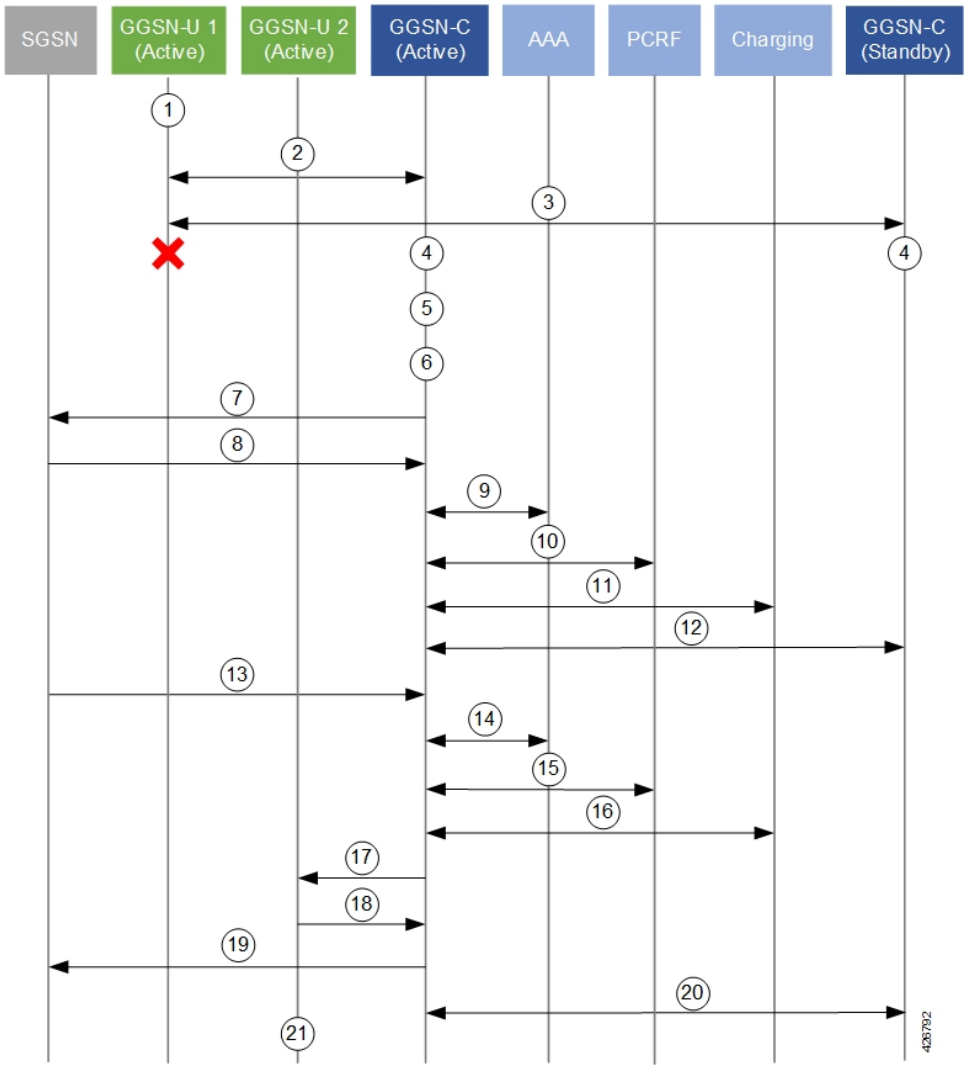

Table 5: GnGp GGSN CP/UP Detach and Re-attach on Path Failure Process

Number	Description
1	UE data sessions are processed by an active GGSN UP.
2	The active GGSN CP monitors GGSN UPs via BFD and Sx-Heartbeat messages.
3	The secondary CP also monitors GGSN UPs via BFD.
4	The active and standby CPs detect a BFD failure on a UP before eNB detection (relays on Sx timers (interval, retransmission, timeout)).
5	The BFD/VPNMGR on the active CP informs the Sx-demux process of a BFDDown event.
6	The Sx-demux process on the active CP initiates a path Failure notice to all Session Managers on the CP.

Number	Description
7	All Session Managers initiate the process of detaching sessions by sending Delete-pdp-context-req messages with no cause code to the SGSN. The detaches are initiated at a pre-defined rate.
8	The SGSN sends Delete-pdp-context-resp messages back to the CP. The SGSN does not page idle UEs with sessions being detached. The SGSN sends E-RAB release messages to active UEs with sessions being detached.
9	The active CP releases the session release with the AAA server(s).
10	The active CP releases the session with the PCRF.
11	The active CP releases the session with the Charging infrastructure.
12	The active CP syncs session detach information with the secondary CP.
13	For UEs re-initiating their session(s), the SGSN sends a Create-pdp-request message to the active CP. The SGSN selects the CP based on load algorithm (DNS, local config etc.).
14	The active CP processes the session attach request with the AAA server(s).
15	The active CP processes the session attach request with the PCRF.
16	The active CP processes the session attach request with the Charging infrastructure.
17	The active CP sends a Sx Session Establishment Request message to an alternate active UP. The CP selects the UP based on its load algorithm.
18	The UP sends a Sx Session Establishment Response message back to the CP.
19	The CP sends a Create-pdp-context response message to the SGSN.
20	The active CP syncs information for the newly attached session with the secondary CP.
21	UE data sessions are now processed by the active GGSN UP.

Additional N+2 Handling Scenarios

Beyond the flows described in the previous sections, the following table provides a description of network function (NF)/system behavior under various conditions with N+2 configured.

Table 6: N+2 Handling Scenarios

ID	Scenario	Handling	Notes
1	Active UP crash	Active CP detects BFD-failure with UP and detaches sessions belonging to that UP. Active CP propagates the disconnects to standby CP through SRP. When UP returns to active, it will re-associate with the active CP.	Detection occurs within the BFD timeout interval. CP Sx monitors BFD.
2	Active CP crash	Active CP switches over to standby CP. Active UP monitors Sx-heartbeat session for both active and standby CPs. Active UP does not purge sessions until ICSR failover time is reached.	Standby CP starts sending Sx-heartbeat upon failover – no sessions are purged by active UP.
3	Standby CP crash	Standby CP comes up and performs checkpoint with active CP to recover sessions	Sessions remain intact on active CP and active UP.
4	Network flaps between active CP and active UP; network between standby CP and active UP remains alive	Active CP detects BFD-Down for UP and initiates session detach processes and disassociates UP. Active CP propagates the disconnects to standby CP through SRP. Active UP monitors Sx-heartbeat with active CP. Active UP waits until configured Sx-heartbeat /path failure detection timeout occurs (>SRP switchover time) before clearing sessions.	

ID	Scenario	Handling	Notes
5	Network flaps between standby CP and active UP; active CP and active UP Sx-heartbeat also down	Active UP detects Sx-path failure. Active UP waits until configured Sx-heartbeat /path failure detection timeout occurs (>SRP switchover time) before clearing sessions. Active CP detects BFD-Down for UP and initiates session detach processes and disassociates UP.	UPs delete the sessions due to Sx-heartbeat timeout.
6	Network flaps between standby CP and active UP; Network between active CP and active UP is alive	Standby CP operates normally. Active CP-active is alive and responds to heartbeat. Active UP operates normally.	
7	Sx is not reachable, however BFD is reachable.	Active UP detects Sx-path failure. Active UP waits until configured Sx-heartbeat/path failure detection timeout occurs (>SRP switchover time) before clearing sessions. Active CP detects Sx-path failure for UP and initiates session detach processes and disassociates UP.	Corner case that is treated as Sx-path failure per current behavior (before N+2).
8	ICSR link between active and standby CPs goes down and standby CP also becomes active (Dual-Active case)	Upon becoming dual-Active, standby CP sends message to active UP with higher metric.	All service IPs advertised by dual-Active standby CP are with higher metric.
9	BGP failure Gn side of active UP	No action is taken in relation to N+2.	
10	BGP failure SGI side of active UP	No action is taken in relation to N+2.	
11	SessMgr crashes on active UP	Session recovery process occurs on active UP.	

ID	Scenario	Handling	Notes
12	Sx-demux crashes on active UP	Sx-demux recovery process occurs on active UP.	
13	VPP crashes on active UP	NPUMgr restarts the UP resulting in BFD loss triggering UP failure detection. Refer to Handling information for IDs 1 and 5 in this table.	
14	VPNMgr crashes on active UP	VPNMgr recovery process occurs on active UP.	
15	BFD crashes on active UP	BFD recovery process occurs on active UP.	
16	Sx-demux crashes on active CP	Sx-demux recovery process occurs on active CP. Sx-demux re-registers for BFD between CP and all UPs as part of recovery and rediscovers the state of each UP. Sx-demux recovers the restart-timestamp from the SessMgr.	It is possible for a UP state change to occur during the Sx-demux recovery on active CP (e.g. UP restarts but still shows as active to CP post recovery). Condition detected as follows: • Sx-demux recovers and CP detects either UP restart timestamp from Sx-heartbeat or UP-failure. • Based on this information, active CP initiates session purging.
17	VPNMgr crashes on active CP	VPNMgr recovery process occurs on active CP. BFDregistration information from recovered from SCT on active CP. Active CP restarts BFD with UP.	

ID	Scenario	Handling	Notes
18	BFD crashes on active CP	BFD recovery process occurs on active CP.	
19	SessMgr crashes on active CP	SessMgr recovery process occurs on active CP.	

Double Failure Handling Scenarios

N+2 double failure scenarios occur when there is a BFD failure followed by another event/failure. The handling of such scenarios is described in the following table.

Table 7: N+2 Double Failure Scenario Handling

ID	Scenario	Handling	Notes
1	Active CP fails while session detaches are in progress	ICSR switchover occurs between CPs. Standby CP becomes active CP. Active CP detects UP failure via BFD. Active CP detects UP restart vis Sx-heartbeat.	Impact: If UP restarts on double failure, it will have no sessions even though the standby CP will have recovered the sessions. These sessions are then cleaned as part of session replacement or session disconnects from UEs. If UP does not restart then the CP-new-active clears the sessions of the failed UP.
2	Standby CP fails while session detaches are in progress	Standby CP checkpoints state information with the Active CP. Information pertaining to deleted sessions is invalidated from active CP.	
3	Active CP determines UP failure due to router flap; Active CP receives UP BFD after initiating session detaching	Once UP BFD down is initially detected, all sessions are detached.	

BFD Flapping and VPC

N+2 uses BFD to monitor the existence/viability of a network path between the session endpoints. By using multihop BFD with loopback endpoints, the BFD session state functions as a proxy for the state of the system to which it connects.

However, a BFD session can go down, or bounce/flap, for reasons other than far-side system failure (e.g. due to ARP storms or router misconfiguration). If the disruption is sufficiently severe and long lasting, it can cause systems on both sides to detect BFD session failure even though both systems are functional.

Configuration adjustments can be made to help offset the occurrence of such events.

The following recommendations are offered based on the platform on which your NFs are deployed:

- VPC-SI: Adjust the BFD multihop-peer settings to increase the BFD detection time to 2-3 sec and the number of retries correspondingly.
- VPC-DI: CF switchover and SF migration can interrupt BFD packet generation and processing for multiple seconds. To prevent BFD session flaps when these events occur, BFD detection time for sessions involving VPC-DI systems must be set to 7 seconds or longer.

Sx-association Scenarios

The following table provides information on associating and disassociating CPs and UPs when using N+2.

Table 8: N+2 Sx-association Scenarios

Scenario	Mechanism(s)
Sx-disassociation from UP to CP	• Sx-demux to disable BFD monitoring with VPNMgr • SAEGW-service is removed • Sx-disassociation from UP
Adding UPs	As part of Day-0: • Add BFD loopback address for UP. • Configurate BFD on CPs. • Add UP Group and configure it for selection on CPs.
Removing UPs	On CP, execute the CLI command to clear subscribers with IP address of UP and keyword to block new sessions being placed on that UP. • Verify that all the subscribers are torn down on UP. • On the UP, execute the CLI command to disassociate from CP. This will disassociate the UP from CP and CP will not choose this UP for further sessions. Verify that all the sessions have been torn down. • On CP, remove the UP from the UP Group. • On CP, execute the CLI command to remove the UP from the UP Group (this will also deregister the BFD monitoring of the UP). • Disable the BFD configurations for monitoring at UP and at CP: no monitor-group CLI command.
UP-initiated Sx-association	Sx-demux on CP starts processing the BFDUp and BFDDown notifications from VPNMGr.

UP-released Sx-association	Sx-demux on CP ignores the BFDUp and BFDDown notifications from VPNMgr.

N+2 and IP Addressing

Loopback IP Addresses

The following is true of BFD loopback addresses in relation to N+2:

- BFD loopback-IP-Address on the active CP and standby CP must be configured on Day-0.
- BFD operates between the active CP and active UP as well as between the standby CP and active UP. As such, all three components must use unique BFD loopback-IP-addresses
- For each CP and UP, configured BFD loopback-IP-addresses must be different from the addresses used for the Sx interfaces, and, in the case of the CPs must also be different from the addresses used for the SRP interface.

IP Address Availability

With the N+2 deployment scenario, UEs may re-attach at a high rate (comparable to the detach rate). To facilitate this process, UPs must have sufficient IP addresses available.

CUPS IP Pool Management includes the capability to provision UPs with "chunks" of addresses. The chunk size and number of pools configured on the CP need to be increased proportionately so as to accommodate the high rate of re-attachments from the CP to UP such that sessions do not get rejected by the UP due to unavailability of IP addresses.

The potential re-attach rate can be roughly estimated by multiplying the number of Session Manager tasks processing UP sessions by 1000 sessions/second.

Address capacity is determined by multiplying the size of the chunk (between 16 and 8192) and the number of IP pools. Both configured on the CP.

Configuring N+2 UP Recovery

To configure N+2 UP Recovery:

1. Configure BFD on the CP and UP.

```
configure
   context bfd_context_name
      ip route static multihop bfd mhbfd_session_name local_endpoint_ip_address
 remote_endpoint_ip_address
      bfd-protocol
         bfd multihop-peer dst_ip_address interval tx_interval min_rx
rx_interval multiplier value
         #exit
      #exit
```

NOTES:

- *bfd_ctx_name* is the name of the context in which BFD is to be configured. This must be the same context in which Sx is configured.
- *mhbfd_session_name* is a name for the BFD session route. Multiple session routes can be created, one for each peer connection.
- *local_endpoint_ip_address* is the IPv4 or IPv6 address corresponding to the local interface in the current context.
- *remote_endpoint_ip_address* is the IPv4 or IPv6 address corresponding to the remote BFD peer.
  - If this route is being configured on the CP, then the remote address is that of the peer UP.
  - If this route is being configured on the UP, then the remote address is that of the peer CP.
- *dst_ip_address* is the IPv4 or IPv6 address corresponding to the remote BFD peer. This must be the same as the *remote_endpoint_ip_address* interface configured for the static multihop BFD route. Multiple peers can be configured, one for each remote peer.
- **interval** *tx_interval* is the transmit interval (in milliseconds) between BFD packets.
- **min_rx** *rx_interval* is the minimum receive interval capability (in milliseconds) between BFD packets.
- **multiplier** *value* the multiplier value used to compute holddown.
- To determine the Detect Time (X), you can use the following calculation:

  Detect Time (X) = **interval** *tx_interval* * **multiplier** *value*

  The recommended value of Detect time (X) is 3 seconds for VPC-SI, and 7 seconds for VPC-DI.

2. Configure the BFD-loopback per context on the CP and UP.

```
configure
   context monitor_ctx_name
      monitor-protocols
         monitor-group monitor_group_name protocol bfd
           session-ctx session_ctx_name local-addr { ipv4_address | ipv6_address
 } remote-address { ipv4_address | ipv6_address }
             #exit
```

NOTES:

- *Monitor_ctx_name* is the name of the context in which BFD monitoring is to be configured. This must be the same context in which Sx is configured.
- *Monitor_group_name* is the name of the group specifying the BFD monitoring parameters. Multiple monitor-groups can be configured.
- *Session_ctx_name* is the name of the context containing the local interfaces over which BFD monitoring will occur. This must be the same context in which Sx is configured.
- **local-addr {** *ipv4_address* **|** *ipv6_address* **}** is the IPv4 or IPv6 address corresponding to the local interface in the specified context.
- **remote-addr {** *ipv4_address* **|** *ipv6_address* **}** is the IPv4 or IPv6 address corresponding to the remote peer with which BFD monitoring will occur.

- If this monitor group is being configured on the CP, then the remote address is that of the UP group.
- If this monitor group is being configured on the UP, then the remote address is that of the CP.

3. Configure the BFD-loopback (remote-IP) within a specific UP-group on the CP:

```
configure
   user-plane-group up_group_name
      peer-node-id { ipv4_address | ipv6_address } monitor-group-name
monitor_group_name
      #exit
```

NOTES:

- *up_group_name* is the name of the UP group containing the data UPs for N+2 UP Recovery will be supported.
  - This cannot be the default group.
  - This group should not contain UPs intended to support IMS/VoLTE.
- **{** *ipv4_address* **|** *ipv6_address* **}** is the IPv4 or IPv6 address of the Sx interface on an active UP that will be part of the UP group. Multiple peer-nodes can be configured within the group. Note that the Sx interface is a different interface from the one that will be used to monitor BFD.
- *monitor_group_name* is the name of the monitoring group the UP will be associated with.

Monitoring and Troubleshooting

Show Commands

show sx peers { full address *peer_ip_address* **| wide }**

show sx peers full address *peer_ip_address*

Displays the Monitor-related information for the specified peer (e.g. VPN context name, group name, and state).

show sx peers wide

Displays "Monitor State" with the default state being "U" for UP, "D" for Down, and "N" for Not Applicable.

show sx-service statistics all

SNMP

The following SNMP traps can be used to monitor N+2 UP Recovery health:

- starBFDSessUp (starentTraps 1276)
- starBFDSessDown (starentTraps 1277)

- starSxPathFailure (starentTraps 1382) – This trap has been updated to include a new cause code: bfd-failure(8)
- starSxPathFailureClear (starentTraps 1383)