

1:1 Redundancy

- Feature Summary and Revision History, on page 1
- Feature Description, on page 2
- How it Works, on page 2
- Configuring 1:1 UPF Redundancy, on page 8
- Monitoring and Troubleshooting, on page 13
- L2 ICSR Support, on page 16

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	5G-UPF
Applicable Platform(s)	VPC-SI SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
Added support for L2 ICSR.	2023.01.0

Revision Details	Release
Support is added for the following functionality: • Zero Accounting Loss in User Plane Function • Early PDU Recovery • Session Prioritization during Recovery • Configuration to change the state of UPF from Pending-Active to Active	2021.02.0
First introduced.	2020.02.0

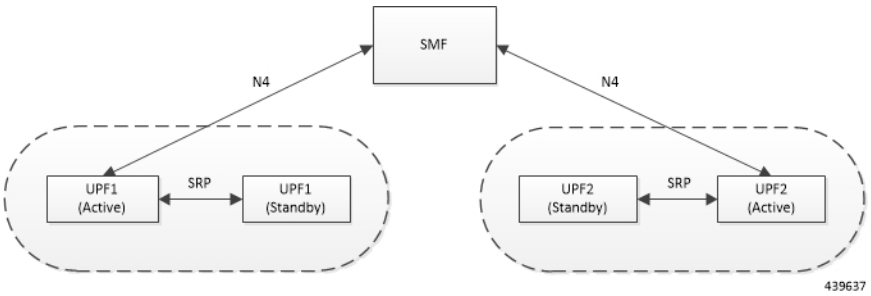
Feature Description

The 1:1 UPF Redundancy feature, for 5G deployment, supports the detection of a failed User Plane Function (UPF) and seamlessly handles the functions of the failed UPF. Each of the Active UPF has a dedicated Standby UPF. The 1:1 UPF Redundancy architecture is based on the UPF to UPF Interchassis Session Recovery (ICSR) connection.

How it Works

The 5G-UPF deployment leverages the ICSR framework infrastructure for checkpointing and switchover of the UPF node as shown in the following figure. The Active UPF communicates to its dedicated Standby UPF through the Service Redundancy Protocol (SRP) link that is provisioned between the UPFs.

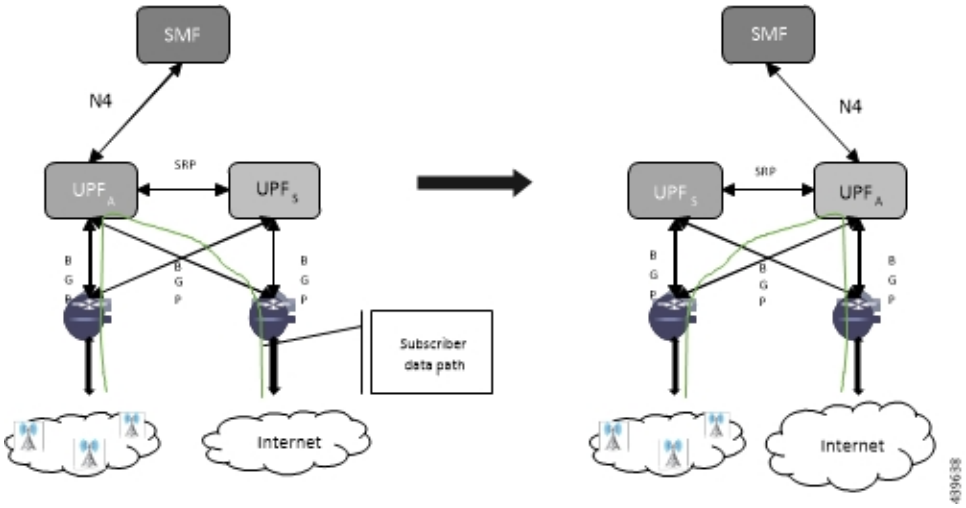
Figure 1: UPF 1:1 Redundancy Using SRP



The Session Management Function (SMF) node does not have the Standby UPF information that is available in the UPF group configuration. Therefore, the SMF is not aware of the UPF redundancy configuration and the switchover event among the UPFs.

The Active UPF communicates to the SMF through the N4 interface address configured in the UPF. The Standby UPF takes over the same Sx/N4 address when it transitions to Active during the switchover event. This implies that the Sx/N4 interface is SRP-activated and is in line with the existing configuration method, therefore UPF switchover is transparent to the SMF.

Figure 2: UPF 1:1 Redundancy Switchover



To make redundancy fully compliant, it addresses the following dependencies on the SRP-based ICSR in the 5G environment.

- Configuration Synchronization (or, Replica Configuration on Standby UPF)
- Sx/N4 Association Checkpoint
- Sx/N4 Link Monitoring

Besides the dependencies listed, the UPF implements data collection and checkpoint procedures specific to the UPF node. For example, checkpointing for IP-pool chunks. The UPF integrates these procedures into the existing ICSR checkpointing framework.

Independent Configuration of Standby UPF

After UPF is up with base configuration (for example, services, contexts, interfaces, and so on), the rest of the configuration (for example, ACS and policy-related configuration) is done through an Ops-center or Redundancy and Configuration Manager (RCM) POD. This configuration is common for both SMF and UPF policies. For SRP redundancy to work, the Active and Standby UPF has same configuration, except SRP-related configuration with which SRP connections are established between the Active and Standby UPF. The RCM configures Active and Standby UPF independently.

BFD Monitor Between Active UP and Standby UP

The Bidirectional Forwarding Detection (BFD) monitors the SRP link between the Active UPF and Standby UPF for a fast failure-detection and switchover. When the Standby UPF detects a BFD failure in this link, it takes over as the Active UPF.

The BFD link can be single-hop or multi-hop.

To configure the BFD monitor, between the Active UP and Standby UP, see *Configuring BFD Monitoring Between Active UPF and Standby UPF*.

Sample Configuration for Multihop BFD Monitoring

Primary UPF:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 209.165.200.225 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.200.225 chassis-to-chassis
      peer-ip-address 209.165.200.225
      bind address 209.165.201.1
    #exit
    interface srp
      ip address 209.165.201.1 255.255.255.224
    #exit
    ip route static multihop bfd bfd1 209.165.201.1 209.165.200.225
    ip route 209.165.201.1 255.255.255.224 209.165.201.1 srp
  #exit
end
```

Backup UPF:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 209.165.201.1 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.201.1 chassis-to-chassis
      peer-ip-address 209.165.201.1
      bind address 209.165.201.1
    #exit
    interface srp
      ip address 209.165.201.1 255.255.255.224
    #exit
    ip route static multihop bfd bfd1 209.165.200.225 209.165.201.1
    ip route 209.165.201.1 255.255.255.224 209.165.200.225 srp
  #exit
End
```

Router between Primary and Backup UPF:

```
config
  context one
    interface one
      ip address 209.165.201.1 255.255.255.224
    #exit
    interface two
      ip address 209.165.200.225 255.255.255.224
    #exit
  #exit
end
```

Sample Configuration for Single-Hop BFD Monitoring

Primary UPF:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.201.1 chassis-to-chassis
      peer-ip-address 209.165.201.1
      bind address 209.165.201.4
    #exit
    interface srp
      ip address 209.165.201.1 255.255.255.224
```

```
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 209.165.201.4
  #exit
end
```

Backup UPF:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.200.225 chassis-to-chassis
      peer-ip-address 209.165.201.4
      bind address 209.165.201.7
    #exit
    interface srp
      ip address 209.165.201.4 255.255.255.224
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 209.165.201.7
  #exit
end
```

VPP Monitor

When SRP VPP monitor is configured, the UPF chassis is SRP Active and if the VPP subsystem fails, then SRP initiates switchover to Standby UPF. Currently, VPP health monitoring is limited to heartbeat mechanism between NPUMgr task and VPP process.

To configure the VPP monitor, see *Configuring VPP Monitor on Active UPF and Standby UPF*.

Sx/N4 Association Checkpoint

Whenever an Active UPF initiates an Sx/N4 association to SMF, the Standby UPF checkpoints this data. This maintains the association information even after the UPF switchover.

The Sx/N4 heartbeat messages are sent and the Active UPF responds back even after back-to-back UPF switchovers.

Sx/N4 Monitor

It is critical to monitor the Sx/N4 interface between the UPF and SMF. The SRP monitoring is enabled on Sx/N4 interface and the existing Sx/N4 heartbeat mechanism is leveraged to detect the monitor failure. The Sx/N4 module on Active UPF, on detecting the failure, informs the SRP VPNMgr to trigger UPF switchover event so that the Standby UPF takes over.



Note Sx/N4 monitoring is available only in the UPF.

It is important to ensure that the SMF Sx/N4 heartbeat timeout is higher than the UPF Sx/N4 heartbeat timeout plus UPF ICSR switchover time. This is to ensure that the SMF does not detect the Sx/N4 path failure during a UPF switchover because of the UPF Sx/N4 monitor failure.

The Standby UPF itself has no independent connectivity to the SMF. The Active UPF Sx/N4 context is replicated to the Standby UPF so that it is ready to takeover during SRP switchover. This implies that when the Active UPF has switched over to Standby because of Sx/N4 monitor failure, the new Standby has no way

of knowing if the UPF to SMF link is working. To prevent a switchback of the new Standby to Active state again due to Sx/N4 monitor failure in new Active, use the **disallow-switchover-on-peer-monitor-fail** keyword in the **monitor sx** CLI command.

After a chassis becomes Standby due to Sx/N4 monitoring failure, the Sx/N4 failure status is not reset even if Sx/N4 up checkpoint is received from the new Active UPF. This is to prevent the new Active to cause an unplanned switchback again due to Sx/N4 monitor failure when the previous cause of switchover itself was Sx/N4 monitor failure. This prevents back-to-back switchovers when SMF is down. The Sx/N4 monitor failure status must be manually reset when the operator is convinced that the network connectivity is normal. To reset, use the new **srp reset-sx-fail** CLI command (see *Resetting Sx/N4 Monitor Failure*) in the Standby chassis.

To configure the Sx/N4 monitor, see *Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF*.

Sx/N4 Monitor—Pending-Active

The UPF chassis can turn into Pending-Active state for one of the following reasons:

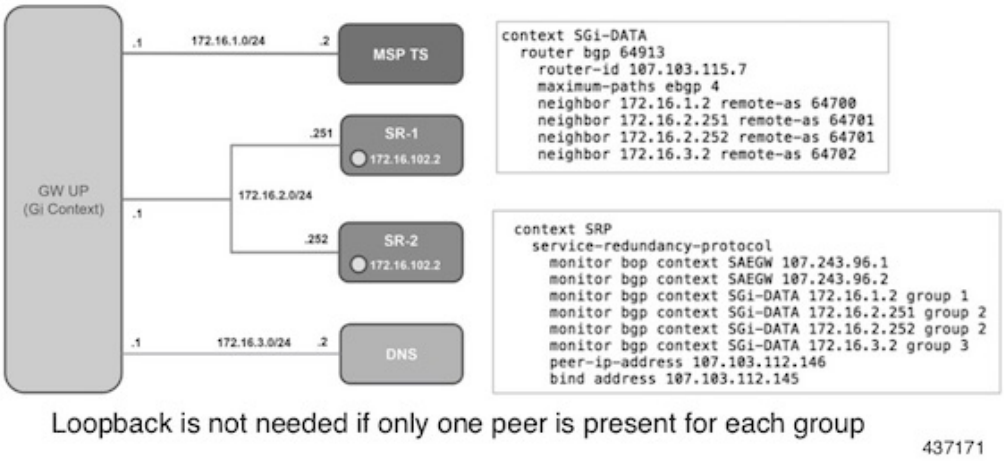
- When Sx/N4 heartbeat times out during SMF upgrade, the Sx/N4 connection is terminated. So, Sx/N4 monitoring failure triggers ICSR switchover in UPF. This switchover causes the old Standby UPF to transition to Pending-Active state. The UPF in Pending-Active state neither receives any Sx/N4 heartbeats from SMF nor any subscriber traffic. As a result, the UPF remains in Pending-Active state indefinitely and can't be utilized without a manual intervention.
- When appropriate procedure to upgrade UPF is not followed, one of the UPF may end up in Pending-Active state. Also, if SMF goes down during the UPF upgrade or if the UPF switchover takes more time than the SMF heartbeat timeout, then one of the UPF remains in Pending-Active state indefinitely.
- When Sx/N4 session times out between SMF and UPF due to network issues, and if a UPF ICSR switchover happens almost simultaneously (Double fault scenario), the UPF in Pending-Active state doesn't transition to active state.

Whenever a UPF chassis turns Pending-Active, start a timer with a callback which forcefully transitions the UPF from Pending-Active to active state. Before forcing the transition, check if the SRP link is up and if the SRP peer is in standby state. If not, restart the timer. The duration of the timer is configurable using **force-pactv-to-actv-timeout** *value_seconds* CLI command (see *Changing UPF state from Pending-Active to Active* section for configuration details). When this CLI command is not configured, the UPF remains in Pending-Active state indefinitely.

BGP Monitor

Configure BGP peer monitor and peer group monitors for the next-hop routers from UPF (both Gi and Gn side). This is the existing ICSR configuration. BGP may run with BFD assist to detect fast BGP peer failure.

Figure 3: BGP Peer Groups and Routing



To configure BGP monitoring and flag BPG monitoring failure, see *Configuring BGP Status Monitoring Between Each UP and Next-Hop Router*.

UPF Session Checkpoints

The Active chassis sends a collection of UPF data as checkpoints to the peer Standby chassis in the following scenarios:

- New call setup
- For every state change in the call
- Periodically for accounting buckets

On receiving these checkpoints, the Standby chassis acts on the data and updates the necessary information either at the call, node, or instance level.

VPN IP Pool Checkpoints

During Sx/N4 Association, the IP pool that is allocated to each of the UPF is sent by SMF to the respective UPF. The VPNMgr receives this message in the UPF and checkpoints the same information to the Standby UPF when the SRP is configured.

The IP pool information is also sent during the SRP VPNMgr restart and during the SRP link down and up scenarios.

Validation of the presence of IP pool information in the Standby is vital before switchover. If the IP pool information is not present, then route advertisement is not possible. Therefore, traffic does not reach the UPF.

External Audit and PFD Configuration Audit Interaction

External Audit management is done in Active UPF. The Session Manager gets a start and complete notification of the Configuration Audit. The Session Manager does not start the External Audit if Configuration Audit is in progress. If the Configuration Audit start-notification arrives when the External Audit is already underway, then the Session Manager raises a flag such that the External Audit restarts when it completes. Restarting the External Audit is necessary because it does not achieve its purpose if it occurs when Configuration Audit is already underway.

Zero Accounting Loss for User Plane Function

Zero accounting loss feature is implemented on the User Plane Function (UPF) so that accounting-data or billing loss is reduced from 18 seconds, which is the default checkpoint time from Active UPF to Standby UPF, or for the configured accounting checkpoint time.

This change in UPF is to support the Gz, Gy, VoGx, and RADIUS URRs. Only planned switchover is supported for zero accounting loss or URR data counters loss. This feature doesn't impact the current ICSR framework or the way checkpointing is done and recovered.

The Sx/N4 usage report is blocked during the “pending active state” until the chassis becomes Active.

Early PDU Recovery for UPF Session Recovery

Early PDU Recovery feature overcomes the earlier limitation of Session Recovery feature wherein it didn't prioritize the CRRs that were selected for recovery. All the CRRs were fetched from the AAAMgr and then the calls were recovered sequentially. The time taken to fetch all the CRRs was a major factor in the perceived delay during session recovery. When a failure occurred, the delay was sometimes long if there were many sessions in a Session Manager. Also, since the calls were recovered in no particular order, the idle sessions were sometimes recovered before active sessions.



Note The Early PDU Recovery feature can recover a maximum of 5-percent sessions.

Session Prioritization during Recovery

Without this functionality, the Session Recovery function didn't prioritize the sessions selected for recovery and loops through all the calls in the call recovery list, and are recovered sequentially when the session recovery is triggered.

As part of Session Prioritization during Recovery, a separate skip list is maintained only for priority calls so that these records can be sent from AAAMgr immediately without going through the loop, thus leading to quicker recovery of the priority calls and reducing the data outage time.

There are two types of sessions at User Plane—Prioritized sessions and normal sessions. Session is considered to be prioritized session based on message priority flag received from SMF and it's recovered first followed by normal calls. These prioritized sessions also take priority in case of early PDU handling. The early PDU of normal calls initiates recovery only when all prioritized sessions are recovered.

In case of critical flush (GR), checkpoints for prioritized sessions are sent first followed by the normal calls. The data of all the calls (both normal and prioritized) are allowed during switchover.



Note The SMF is responsible to set the priority flags for all the calls. The UPF uses the priority call details that are received from the SMF for the Session Prioritization feature.

Configuring 1:1 UPF Redundancy

The following sections provide information about the CLI commands available in support of the feature.

Configuring BFD Monitoring Between Active UPF and Standby UPF

Use the following configuration to configure Bidirectional Forwarding Detection (BFD) monitoring on the Active UPF and Standby UPF. Configure this command in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
         [ no ] monitor bfd context context_name { ipv4_address | ipv6_address }
{ chassis-to-chassis | chassis-to-router }
         end
```

NOTES:

- **no**: Disables BFD monitoring on the Active and Standby UPF.
- **context** *context_name* : Specifies the context that is used. It refers to the context where the BFD peer is configured (SRP context).

  context_name must be an existing context expressed as an alphanumeric string of 1 through 79 characters.
- *ipv4 _address* | *ipv6_address*: Defines the IP address of the BFD neighbor to be monitored, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

  It refers to the IP address of the configured BFD (ICSR) peer.
- **chassis-to-chassis** | **chassis-to-router**:

  chassis-to-chassis: BFD runs between primary and backup chassis on non-SRP links.

  chassis-to-router: BFD runs between chassis and router.



Caution Don't use the **chassis-to-router** keyword for BFD monitoring on the SRP link between the Active UPF and the Standby UPF.

- This command is disabled by default.

Configuring BGP Status Monitoring Between Each UPF and Next-Hop Router

Use the following commands to configure Border Gateway Protocol (BGP) monitoring between each UPF and next-hop router. The command is configured in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
         [ no ] monitor bgp context bgp-session-context-name [
nexthop-router-ipv4-address | nexthop-router-ipv6-address ] { vrf
bgp-session-vrf-name } { group group-number  }
         end
```

NOTES:

- **no**: Disables BGP status monitoring on the UPF.

- **bgp context** *bgp-session-context-name*: Specifies the context where BGP peer is configured. *bgp-session-context-name* specifies the context string.
- **nexthop-router-ipv4-address | nexthop-router-ipv6-address**: Specifies the configured BGP peer IPv4 or IPv6 address to monitor.
- **vrf** *bgp-session-vrf-name*: Specifies the BGP VPN Routing and Forwarding (VRF) instance. *bgp-session-vrf-name* specifies the VRF name.
- **group** *group-number* : Specifies the BGP peer group where the BGP peer should be included. *group-number* specifies the group number.

  On implementing this keyword, the behavior is as follows:

  - If any BGP peer in that group is up, the BGP peer group is up.

    Omitting group configuration for a BGP monitor includes that monitor in group 0.

    BGP group 0 monitors in a context from an implicit group. Each context forms a separate BGP group 0 implicit monitor group.

    If any BGP peer group is down, BGP monitor is down.

- This command is disabled by default.

Alternate Algorithm to Flag BGP monitoring failure

In this release, an alternate (new) algorithm is introduced to flag BGP monitoring failure.

Use the following commands to flag BGP monitor failure on a single BGP peer (User Plane Function) failure. This command is configured in the SRP Configuration Mode.

```
configure
   context context_name
      service-redundancy-protocol
         [ no ] monitor bgp exclusive-failover
         end
```

NOTES:

- **no**: Disables flagging of BGP monitor failure on a single BGP peer failure.
- On implementing the new **exclusive-failover** keyword, the behavior is as follows:
  - BGP peer group is Up if any BGP peer in that group is Up.
  - Including a BGP peer in group 0 is same as making it non-group (omitting group).
  - BGP monitor is down if any BGP peer group or any non-group BGP peer is down.
- This command is disabled by default.

Configuring Sx/N4 Monitoring on the Active UPF and Standby UPF

Use the following configuration to configure Sx/N4 monitoring on the Active UPF and Standby UPF. This command is configured in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      [ no ] monitor sx [ { context context_name | bind-address {ipv4_address
|ipv6_address } | { peer-address { ipv4_address | ipv6_address } } ]
      end
```

NOTES:

- **no**: Disables Sx/N4 monitoring on the Active and Standby UPF.
- **context** *context_name* : Specifies the context of the Sx/N4 service.

  context_name must be an existing context expressed as an alphanumeric string of 1 through 79 characters.
- **bind-address {** *ipv4 _address* **|** *ipv6_address* **}**: Defines the service IP address of the Sx/N4 service, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

Note The IP address family of the **bind-address** and **peer-address** must be same.

- **peer-address {** *ipv4 _address* **|** *ipv6_address* **}**: Defines the IP address of the Sx/N4 peer, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.
- **disallow-switchover-on-peer-monitor-fail**:

  Prevents the switchback of the UPF to Active state when the working status of the UPF to SMF link is unknown.
- It's possible to implement this CLI command multiple times for monitoring multiple N4/Sx connections.
- The Sx/N4 monitor state goes down when any of the monitored Sx/N4 connections are down.
- This command is disabled by default.
- A maximum of 18 peer nodes can be configured.

Configuring VPP Monitor on Active UPF and Standby UPF

Use the following configuration to configure Vector Packet Processing (VPP) monitor to trigger UPF switchover on the Active UPF if VPP goes down. This command is configured in the SRP Configuration Mode.

```
configure
  context context_name
    service-redundancy-protocol
      monitor system vpp delay-period seconds
      end
```

NOTES:

- If previously configured, use the **no monitor system vpp** CLI command to disable VPP monitoring on the Active and Standby UPF.
- **vpp delay-period** *seconds* : Specifies the delay period in seconds for a switchover, after a VPP failure. *seconds* must be in the range of 0 through 300.

If the delay period is a value greater than zero (0), then the switchover is initiated after the specified delay period when VPP fails. The last VPP status notification within the delay period is the final trigger for switchover action. The default value is 0 seconds, which initiates an immediate switchover.

The need for delay is to address the scenario wherein the VPP is temporarily down and the revival is in process. This implies that a switchover may not be necessary.

- This command is disabled by default.

Preventing User Plane Function Switchback

Use the following configuration to prevent the switchback of the new Standby UPF to Active state again due to Sx/N4 monitor failure in the new Active.

```
configure
  context context_name
    service-redundancy-protocol
      monitor sx disallow-switchover-on-peer-monitor-fail timeout seconds

       end
```

Use either of the following CLIs to allow switchback of the new Standby UPF to Active state.

```
no monitor sx disallow-switchover-on-peer-monitor-fail
```

Or

```
monitor sx disallow-switchover-on-peer-monitor-fail timeout 0
```

NOTES:

- **no**: Disables prevention of switchover.
- **disallow-switchover-on-peer-monitor-fail [timeout** *seconds* **]** : Prevents the switchback of the UPF to Active state when the working status of the UPF to SMF link is unknown.

  timeout *seconds*: Timeout after which the switchback is allowed even if the Sx/N4 failure status is not reset in the Standby peer. The valid values range from 0 through 2073600 (24 days).



Note Assigning 0 seconds as the timeout allows unplanned switchover.

If **timeout** keyword is not specified, the Active chassis waits indefinitely for the Sx/N4 failure status to be reset in the Standby peer.

- The default configuration is to allow unplanned switchover due to Sx/N4 monitor failure in all conditions.



Note Manual planned switchover is allowed irrespective of whether this CLI is configured or not.

Preventing Dual Active Error Scenarios

Use the following CLI configuration in CP to prevent dual Active error scenarios for UPF 1:1 redundancy.

```
configure
    user-plane-group group_name
       sx-reassociation disabled
       end
```

NOTE:

- **sx-reassociation disabled**: Disables UP Sx reassociation when the association already exists with the CP.

Resetting Sx/N4 Monitor Failure

Use the following configuration only on the Standby chassis to reset the Service Redundancy Protocol (SRP) Sx/N4 monitor failure information. This command is configured in the Exec Mode.

```
srp reset-sx-fail
```

Changing UPF State from Pending-Active to Active

Use the following configuration to change the UPF chassis state from Pending-Active to Active.

```
configure
  context context_name
    service-redundancy-protocol
      force-pactv-to-actv-timeout value_seconds
```

NOTES:

- *value_seconds*: Specifies the timeout value in seconds and must be in the range of 1-300.
- Use the **show config context** *context_name* CLI command to verify the configuration.

Monitoring and Troubleshooting

This section provides information regarding the CLI command available in support of monitoring and troubleshooting the feature.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of this feature.

show srp monitor bfd

The output of this CLI command contains the following fields for the 5G UPF 1:1 Redundancy feature:

- Type:
  - (A) - Auth. probe

  - (B) - BGP
  - (D) - Diameter
  - (F) - BFD
  - (E) - EGQC
  - (C) - Card
  - (V) - VPP
- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down
- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp monitor bgp

The output of this CLI command contains the following fields for the 5G UPF 1:1 UPF Redundancy feature:

- Type:
  - (A) - Auth. probe
  - (B) - BGP
  - (D) - Diameter
  - (F) - BFD
  - (E) - EGQC
  - (C) - Card
  - (V) - VPP
  - (S) - Sx
- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down

- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp monitor sx

The output of this CLI command contains the following fields in support of Sx/N4 monitor status:

- Type:
  - (A) - Auth. probe
  - (B) - BGP
  - (D) - Diameter
  - (F) - BFD
  - (E) - EGQC
  - (C) - Card
  - (V) – VPP
  - (S) - SX
- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down
- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp monitor vpp

The output of this CLI command contains the following fields for the 5G UPF 1:1 UPF Redundancy feature:

- Type:
  - (A) - Auth. probe
  - (B) - BGP

- (D) - Diameter
- (F) - BFD
- (E) - EGQC
- (C) - Card
- (V) - VPP

- State:
  - (I) - Initializing
  - (U) - Up
  - (D) - Down

- GroupId
- IP Addr
- Port
- Context (VRF Name)
- Last Update

show srp statistics

The output of this CLI command contains the following fields for the Sx/N4 Monitor—Pending-Active functionality:

- Pending-active timer started
- Pending-active timer stopped
- Pending-active to Active forced
- Pending-active to Active force-failed
- Pending-active to Active force-skipped - peer-not-sby
- Pending-active to Active force-skipped - not-PActv

L2 ICSR Support

Feature Description

UPF supports Layer 2-based (L2) ICSR in a two-server architecture to disambiguate packet forwarding between Active and Standby UPFs. UPF uses Gratuitous Address Resolution Protocol (GARP) and Unsolicited Neighbor Advertisement messages to achieve L2 ICSR.

L2 ICSR between SRP peers is configurable in the SRP context. UPF also configures a loopback interface under the same subnet as that of the physical interface. UPF supports SRP BFD monitoring to avoid the Active to Active scenario.

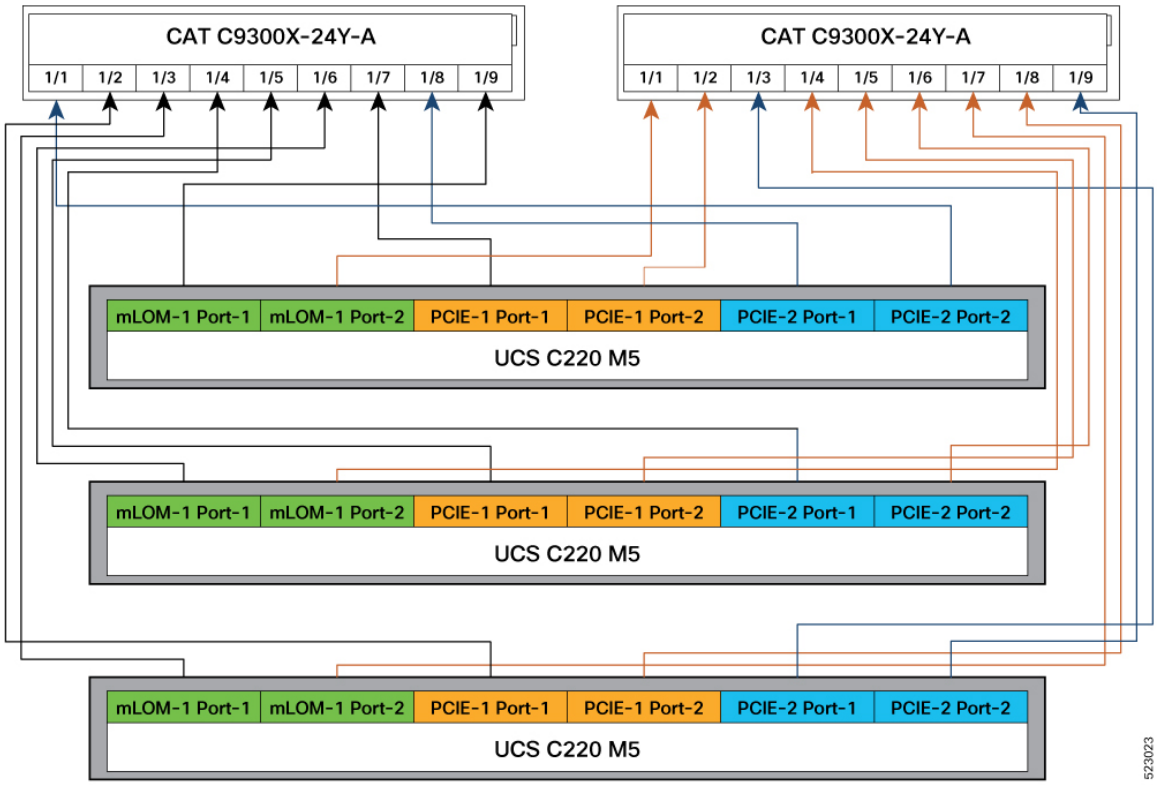
SRP Loopback Interface

SRP loopback interface is the loopback interface for Service IP (N3) or Next-hop IP for IP-Pools (N6). The IP addresses of these loopback interfaces must be on the same subnet as the ethernet interfaces of the context. The SRP loopback interfaces have the same IP address on both Active and Standby UPFs.

- When UPF is Standby, SRP loopback interfaces will not send ARP reply or GARP packets. For IPv6, the UPF does not respond with the Solicited or the Unsolicited Neighbor Advertisement message.
- When UPF is Standby, SRP loopback interfaces are enabled with an IP address attached inside boxer so that the UPF services bound to these interfaces remain in running state.
- When UPF transitions from any other state to the Active state, the SRP loopback interfaces send out GARP packets. For IPv6, the UPF sends Unsolicited Neighbor Advertisement messages.

The following figure illustrates the 3RU wiring diagram for non-bonded interfaces.

Figure 4: 3RU Wiring for Non-Bonded Interfaces



How it Works

UPF directly connects gNB on N3 router or Server on N6 router in the same subnet. UPF ICSR will not be visible to gNB/N3 and Server/N6.

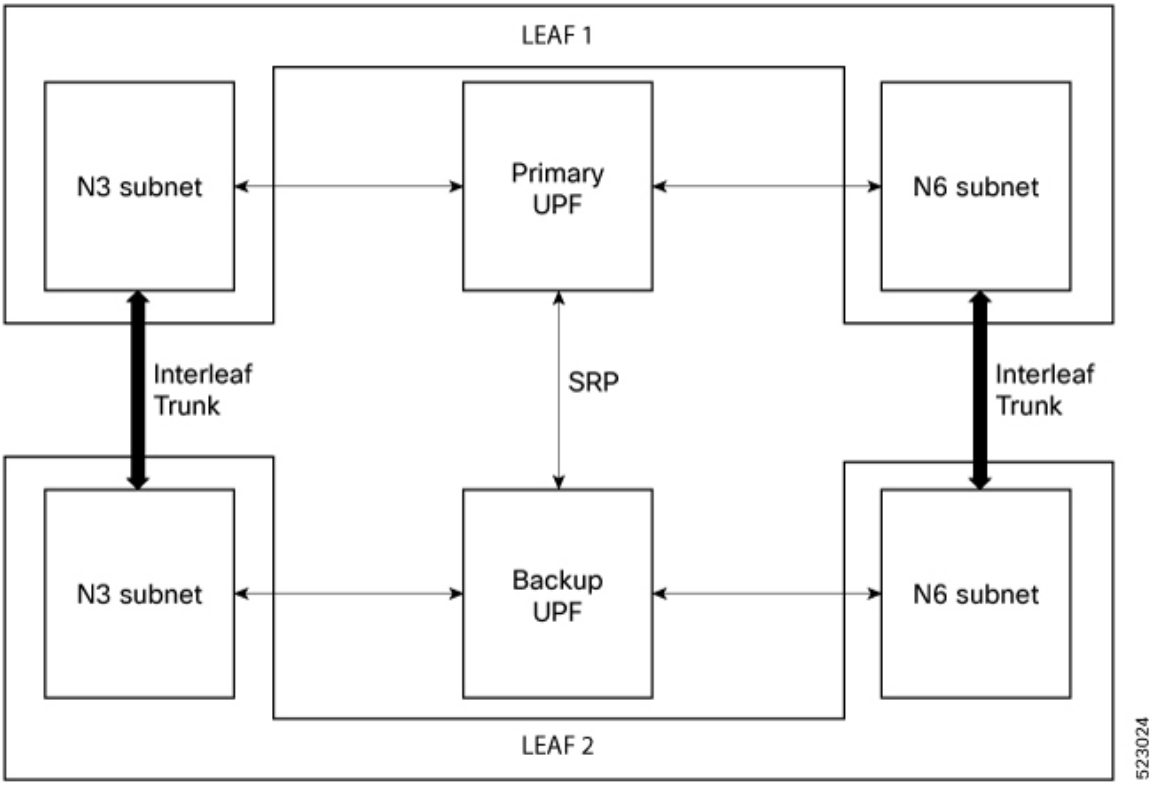
The L2 switch supports BFD. On all N3 and N6 subnets, the switch must have an IP address.

UPF requires Chassis-Chassis (C-R) BFD and Chassis-Router (C-R) BFD for the L2 ICSR solution to work. To avoid the Active-Active state, BFD runs between C-C and C-R. C-C BFD must have a higher timeout value than C-R BFD.

For a two-leaf network design and non-bonded UPF interfaces, each UPF connects to a separate leaf. The IP addresses of the L2 switch are part of the leaf.

The following figure illustrates the high-level diagram of L2 switch wiring for non-bonded interfaces.

Figure 5: L2 Switch Wiring for Non-Bonded Interfaces



State Machine

This table provides information about the UPF states.

Table 1: States Machine

S.No.	Previous State	Internal Fault	C-R BFD	C-C BFD	SRP Link State	Next State	Pull Down BFD
1	X	Yes	X	X	X	Standby	Yes
2	X	X	Any-Down	X	X	Standby	Yes
3	X	No	All-Up	Any-Down	X	Active	No

S.No.	Previous State	Internal Fault	C-R BFD	C-C BFD	SRP Link State	Next State	Pull Down BFD
4	X	No	All-Up	All-Up	Down	No-Change	No
5	X	No	All-Up	All-Up	Up	SRP-Protocol -Decided	No

NOTES:

- For C-R BFDs going from "Not-All-Up" to "All-Up", a delay is introduced in SRP State processing to allow all C-C BFDs to reach steady state.
- To avoid the dual-active scenario, ensure that the states in row 3 of the above table do not occur in both UPFs. This must not happen in a properly functioning and configured system.

Configuring L2 ICSR in SRP Context

To enable or disable L2 ICSR between SRP peers, use the following configuration:

```
configure
  context context_name
    service-redundancy-protocol
      l2-icsr [ disable | enable ]
      end
```

NOTES:

- **disable**: Disable L2 ICSR between SRP peers. This is the default setting.
- **enable**: Enable L2 ICSR between SRP peers.

Configuring C-C and C-R BFD from Primary and Backup Chassis

To configure the C-C and C-R BFD from a Primary and Backup Chassis use the following configuration. The C-C BFD configures as multi-hop to give a different timeout value than the C-R BFD.

```
context SAEGW
bfd-protocol
  bfd multihop-peer 209.165.200.227 interval 50 min_rx 50 multiplier 6
#exit
interface S5_SGW_PGW
  ip address 209.165.200.226 255.255.255.0
  ipv6 address 2001:DB8:A:B::1/64 secondary
  bfd interval 50 min_rx 50 multiplier 3
#exit
ip route static bfd S5_SGW_PGW 209.165.200.225
ip route static multihop bfd bfdcc 209.165.200.226 209.165.200.227
ip igmp profile default
#exit
```

Verifying L2 ICSR Configuration

To verify the configuration, use the following commands:

- **show config context srp**

- **show srp monitor all**
- **show srp info**

The following is a sample output of **show config context srp** command:

```
[local]upf# show config context srp
    context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      checkpoint session duration non-ims-session 0
      checkpoint session duration ims-session 0
      force-pactv-to-actv-timeout 1
      l2-icsr enable
      [output truncated]
```

The following is a sample output of **show srp monitor all** command:

```
[local]upf# show srp monitor all
+----- Type: (A) - Auth. probe (B) - BGP (D) - Diameter (F) - BFD (E) - EGQC (C) - Card
| (V) - VPP (S) - Sx
|
|+---- State: (I) - Initializing (U) - Up (D) - Down
||
||+--- GroupId
|||
vvv IP Addr Port Context(VRF Name) Last Update
-------------------------------------------------------------------------------
FU0 209.165.200.225 3785 SAEGW Thu Jan 05 01:13:08 2023
FU0 209.165.200.227 3785 SAEGW Thu Jan 05 01:13:08 2023
FU0 209.165.200.230 3785 ISP Thu Jan 05 01:13:08 2023
FU0 209.165.200.231 3785 ISP Thu Jan 05 01:13:08 2023
-------------------------------------------------------------------------------
[local]upf#
[local]upf# context SAEGW
[SAEGW]upf# show bfd neighbors >>>>>
All the time intervals are in ms
OurAddr NeighAddr LD/RD DtctTime State Intf
209.165.200.226 209.165.200.225 1/1090519046 150 Up S5_SGW_PGW
209.165.200.226 209.165.200.227 2/1 300 Up
Total Sessions 2
Sessions in UP state 2
```

The following is a sample output of **show srp info** command:

```
sample o/p:
[local]l10-ru34-up1# show srp info
Service Redundancy Protocol:
-------------------------------------------------------------------------------
Context: srp
Local Address: 209.165.201.234
Chassis State: Active
Chassis Mode: Primary
Chassis Priority: 47299
Local Tiebreaker: 00-50-56-AE-EE-8D
Route-Modifier: 33

Peer Remote Address: 209.165.201.235
Peer State: Standby
Peer Mode: Backup
Peer Priority: 1673325643
Peer Tiebreaker: 00-00-00-00-00-00
Peer Route-Modifier: 0
Last Hello Message received: Tue Jan 10 04:40:43 2023 (2 seconds ago)
```

```
Peer Configuration Validation: Complete
Last Peer Configuration Error: None
Last Peer Configuration Event: Tue Jan 10 04:32:31 2023 (494 seconds ago)
Last Validate Switchover Status: None
Connection State: Connected
Next Peer Audit Scheduled:
Periodicity Audit Scheduled : -
Daily-Start-time Audit Scheduled : -
Cron-Job Audit Scheduled : -
Peer Audit State: Unknown
Last Peer Audit Type: Unknown
Last Peer Audit Successful: -
DSCP Configuration:
Control flow : be
Session flow : be
Feature Configured-status Operational-status
allow-volte-data-traffic: Disabled Enabled
allow-all-data-traffic: Disabled Enabled

[local]l10-ru34-up1#
```

OAM Support

This section describes operations, administration, and maintenance information for this feature.

Show Command Support

Use the **show ip interface** or **show ipv6 interface** command to display the GARP statistics.

The following is a sample output of this command:

```
[EPC2]l10-ru32-up1# show ip interface
Interface Name                    : S5_SGW_PGW
Interface Type                    : Broadcast
Description                       :
VRF                               : None
IP State                          : UP (Bound to 1/10 vlan id 2141, 802.1P prior 0, ifIndex
 17432595)
IP Address                        : 209.165.200.226
Subnet Mask                       : 255.255.255.0
Bcast Address                     : 209.165.200.255
MTU                               : 1500
Resoln Type                       : ARP
ARP timeout                       : 1200 secs
L3 monitor LC-port switchover     : Disabled
Number of Secondary Addresses     : 1
    IPv6 Address: 2001:DB8:A:B::1/64
GARP sent through this interface : 45
GARP sent for this interface     : 0
GARP sent for/through this interface successfully: 45
Last GARP sent for/through this interface: Tue Dec 06 18:40:26 2022
```