

1:1 User Plane Redundancy for 4G CUPS

- Revision History, on page 1
- Feature Description, on page 1
- How it Works, on page 1
- Configuring 1:1 User Plane Redundancy for 4G CUPS, on page 11
- Monitoring and Troubleshooting, on page 16

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

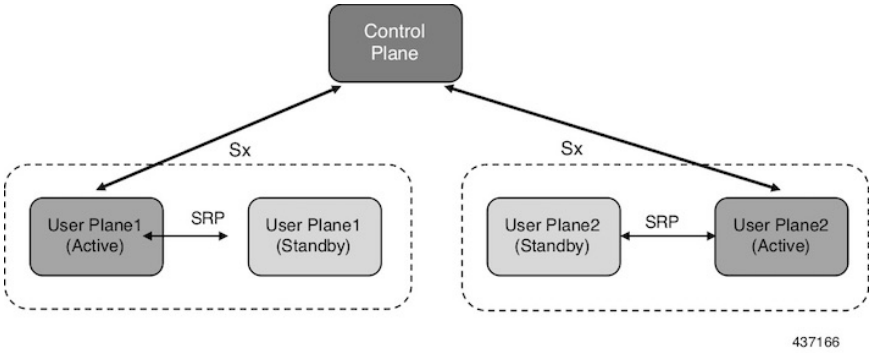
The 1:1 User Plane Redundancy for 4G CUPS feature supports the detection of a failed User Plane (UP) and handles seamlessly the functions of the failed UP. Each of the Active UPs has a dedicated Standby UP. The 1:1 UP redundancy architecture is based on the UP to UP Interchassis Session Recovery (ICSR) connection.

How it Works

This section briefly describes how 1:1 User Plane Redundancy for 4G CUPS feature works.

The 4G CUPS deployment leverages the ICSR framework infrastructure for checkpointing and switchover of the UP node as shown in the following figure. The Active UP communicates to its dedicated Standby UP via the Service Redundancy Protocol (SRP) link that is provisioned between the UPs.

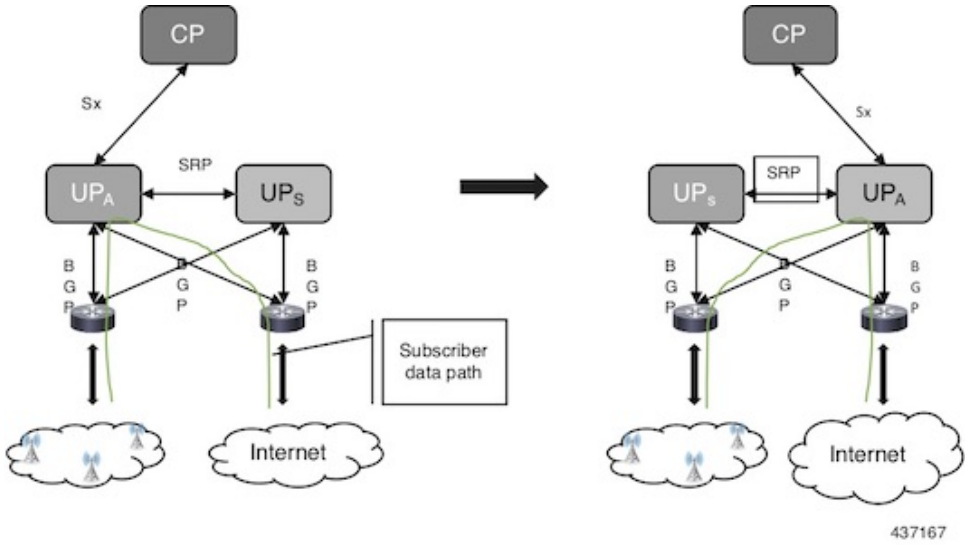
Figure 1: UP 1:1 Redundancy Using SRP



The Control Plane (CP) node does not have the Standby UP information that is available in the UP group configuration. Therefore, the CP is not aware of the UP redundancy configuration and the switchover event among the UPs.

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

Figure 2: UP 1:1 Redundancy Switchover



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

- Synchronization of PFD Configuration
- Sx Association Checkpoint
- Sx Link Monitoring

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

Figure 3: CP-CP ICSR with 1:1 UP Redundancy, before CP Switchover

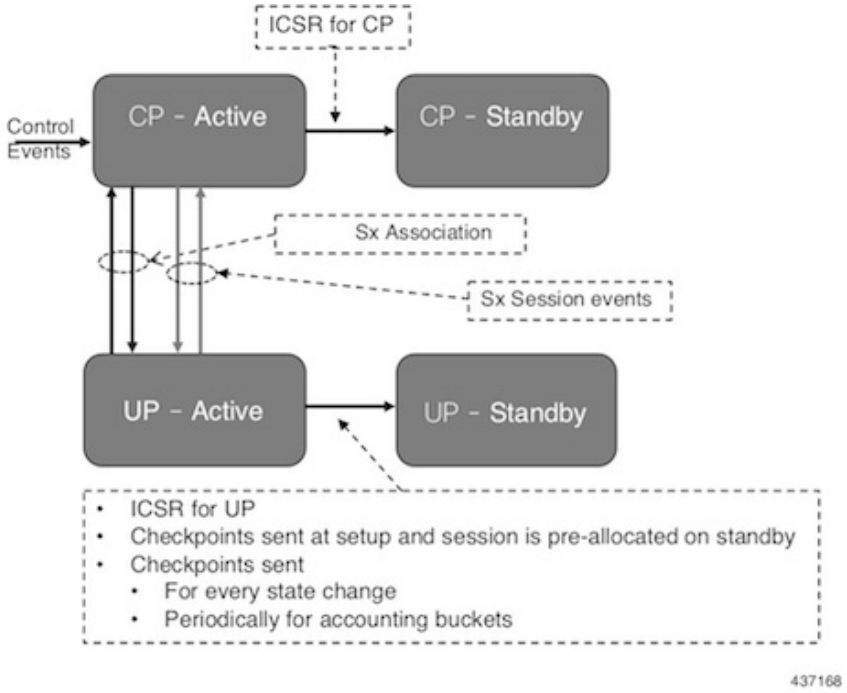
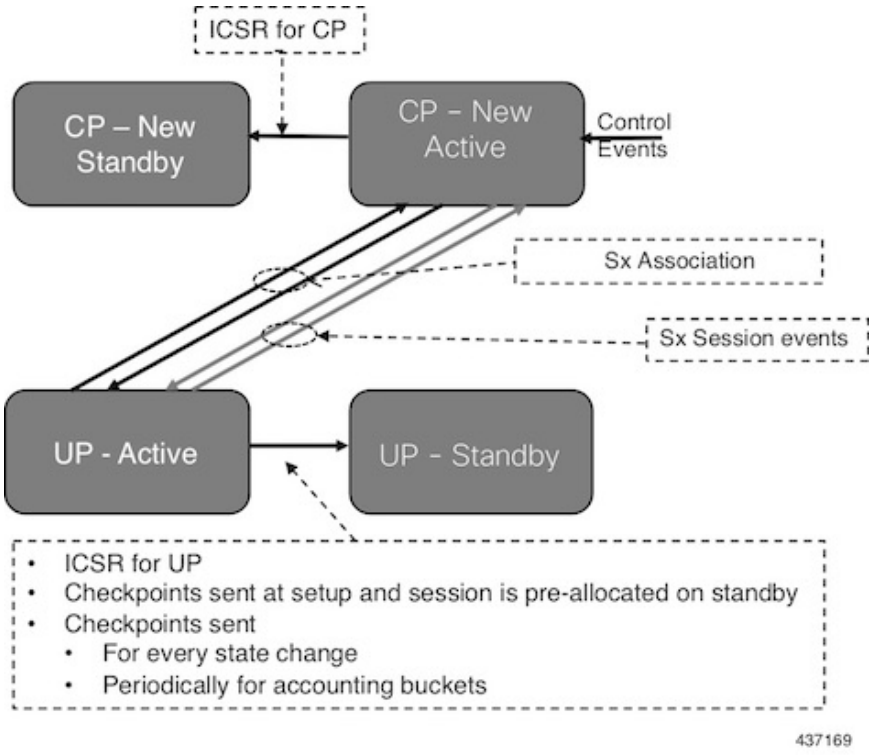


Figure 4: CP-CP ICSR with 1:1 UP Redundancy, After CP Switchover

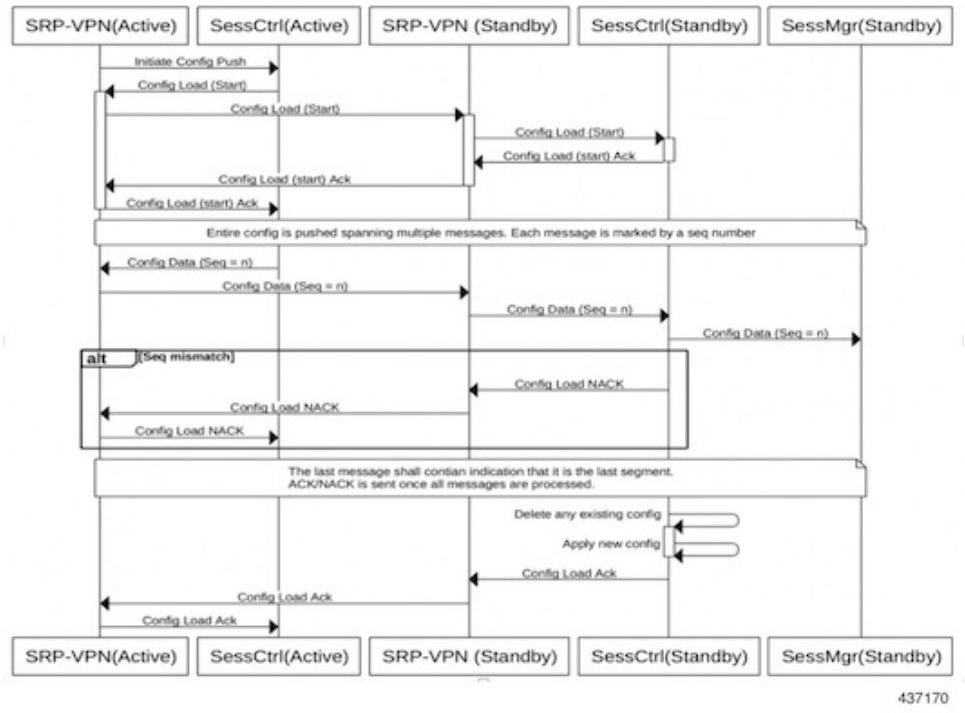


Synchronization of PFD Configuration

The CP node pushes the UP configuration via the Packet Flow Description (PFD) messages. The CP sends the PFD configuration from the Active UP to the Standby UP because the Sx IP address of the UP is SRP-activated over the Active UP and Standy UP.

The SRP VPN Manager provides the transport between UPs and the Session Controller in the Active UP anchors the configuration push. The following illustration lists the sequence of events.

Figure 5: Synchronizing PFD Configuration



BFD Monitor Between Active UP and Standby UP

The BFD monitors the SRP link between the Active UP and Standby UP for a fast failure detection and switchover. When the Standby UP detects a BFD failure in this link, it takes over as the Active UP.

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

Note The recommendation is that the SRP bind interface must be an Ethernet interface that attaches to the card service Port. In a loopback address, the recommendation is to ensure that the BFD control packets traverse only though one service port. If it is the ECMP, ensure that the route convergence time does not exceed the BFD timeout.

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

Sample Configuration for Multihop BFD Monitoring

Primary UP:

```
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 209.165.200.227
```

```
    #exit
    interface srp
      ip address 209.165.200.227 255.255.255.224
    #exit
    ip route static multihop bfd bfd1 209.165.200.227 209.165.200.225
    ip route 192.168.210.0 255.255.255.224 209.165.200.228 srp
  #exit
end
```

Backup UP:

```
config
  context srp
    bfd-protocol
      bfd multihop-peer 209.165.200.227 interval 50 min_rx 50 multiplier 20
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.200.227 chassis-to-chassis
      peer-ip-address 209.165.200.227
      bind address 209.165.200.225
    #exit
    interface srp
      ip address 209.165.200.225 255.255.255.224
    #exit
    ip route static multihop bfd bfd1 209.165.200.225 209.165.200.227
    ip route 192.168.209.0 255.255.255.224 209.165.200.226 srp
  #exit
End
```

Router between Primary UP and backup UP:

```
config
  context one
    interface one
      ip address 209.165.200.228 255.255.255.224
    #exit
    interface two
      ip address 209.165.200.226 255.255.255.224
    #exit
  #exit
end
```

Sample Configuration for Single Hop BFD Monitoring

Primary UP:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 255.255.255.230 chassis-to-chassis
      peer-ip-address 255.255.255.230
      bind address 209.165.200.227
    #exit
    interface srp
      ip address 209.165.200.227 255.255.255.224
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 255.255.255.230
  #exit
end
```

Backup UP:

```
config
  context srp
    bfd-protocol
    #exit
    service-redundancy-protocol
      monitor bfd context srp 209.165.200.227 chassis-to-chassis
      peer-ip-address 209.165.200.227
      bind address 255.255.255.230
    #exit
    interface srp
      ip address 255.255.255.230 255.255.255.224
      bfd interval 50 min_rx 50 multiplier 10
    #exit
    ip route static bfd srp 209.165.200.227
  #exit
end
```

VPP Monitor

The SRP VPP monitor initiates a switchover to Standby UP when the VPP subsystem fails.

Note The VPP monitor is available only on the VPC-SI instance UP. It is not available in the hybrid CUPS ASR 5500 UP because the card level redundancy handles the VPP failure on the ASR 5500. If VPP causes multiple card failures, then SRP card monitor must be used.

To configure the VPP monitor, see "Configuring VPP Monitor on Active UP and Standby UP."

Sx Association Checkpoint

Whenever an Active UP initiates a Sx association to the configured CP node, the Standby UP checkpoints this data. This maintains the association information even after the UP switchover.

The Sx heartbeat messages sends and the Active UP must responds even after back-to-back UP switchovers.

Sx Monitor

It is critical to monitor the Sx interface between the UP and CP. Enabling the Sx heartbeat functionality is essential because it helps detect a monitor failure.

Note Sx monitoring is available only in the UP.

The Sx interface on the Active UP detects failure and informs the SRP VPN Manager to trigger the UP switchover event such that the Standby UP takes over.

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

Preventing Control Plane Heartbeat Time Out

There is a minor possibility that the CP heartbeat times out during the UP ICSR switchover. Follow these steps to mitigate it:

1. Remove the Sx heartbeats from the CP toward the UPs.

2. If the former is not possible, then ensure that the Sx heartbeats from the CP toward the UP have multiple retry timeout. Also ensure that the number of retries is greater than the UP Sx heartbeat timeout plus UP ICSR switchover time.

   For example:

   A = CP heartbeat interval (*sx-protocol heartbeat interval*)

   B = CP heartbeat max retransmissions (*sx-protocol heartbeat max-retransmissions*)

   C = CP heartbeat retransmission timeout (*sx-protocol heartbeat retransmission-timeout*)

   D = UP heartbeat interval (*sx-protocol heartbeat interval*)

   E = UP heartbeat max retransmissions (*sx-protocol heartbeat max-retransmissions*)

   F = UP heartbeat retransmission timeout (*sx-protocol heartbeat max-retransmissions*)

   G = Switchover time (including BGP route convergence time)

   Therefore, the formula for successful Sx monitor failure switchover is:

   B * C > D + (E * F) + G

   Example Values:

   CP:

   A:

   sx-protocol heartbeat interval 60

   B:

   sx-protocol heartbeat max-retransmissions 10

   C:

   sx-protocol heartbeat retransmission-timeout 10

   UP:

   D:

   sx-protocol heartbeat interval 30

   E:

   sx-protocol heartbeat max-retransmissions 3

   F:

   sx-protocol heartbeat retransmission-timeout 3

   BGP:

   G: Example route converge time = 30 sec

   Therefore, B * C > D + (E * F) + G

   => 10 * 10 > 30 + (3 * 3) + 30

   => 100 > 69

   A maximum value of B is 15 and max value of C is 20. Therefore, configure the Sx monitor failure detection and UP switchover (D + (E * F) + G) to withstand a maximum delay of 15 * 20 = 300 sec, that is, 5 min.

To minimize the BGP route convergence time (G), run the BGP with BFD fail-over.

To configure the Sx monitor, see "Configuring Sx Monitoring on the Active UP and Standby UP."

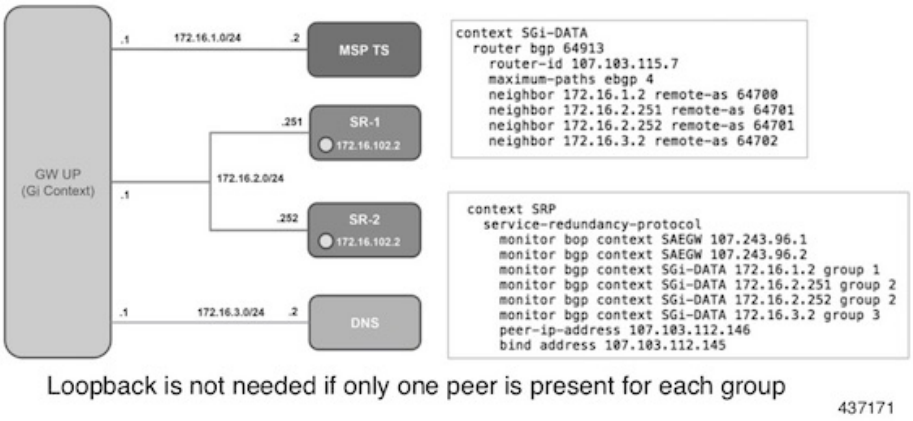
The Standby UP itself has no independent connectivity to the CP. The Active UP Sx context is replicated to the Standby UP such that it is ready to takeover during SRP switchover. This implies that when the Active UP has switched over to Standby because of Sx monitor failure, the new Standby has no way of knowing if the UP to CP link is working. To prevent a switchback of the new Standby to Active state again due to Sx monitor failure in new Active, use the **disallow-switchover-on-peer-monitor-fail** keyword in the new **monitor sx** CLI command.

After a chassis becomes Standby due to Sx monitoring failure, the Sx failure status is not reset even if Sx up checkpoint is received from the new Active UP. This is to prevent the new Active to cause an unplanned switchback again due to Sx monitor failure when the previous cause of switchover itself was Sx monitor failure. This prevents back-to-back ping-pong type of switchovers when CP is down. The Sx 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 Monitor Failure") in the Standby chassis.

BGP Monitor

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

Figure 6: BGP Peer Groups and Routing



To configure BGP monitoring and flag BPG monitoring failure, see Flagging BGP Monitoring Failure, on page 12.

UP Session Checkpoints

The Active chassis sends a collection of UP 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 level or node or instance level.

VPN IP Pool Checkpoints

Along with the PFD configuration message, the CP sends the IP pool allocation to each UP. The VPN manager receives this message in the UP and checkpoints the same information to the Standby UP 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 UP.

External Audit and PFD Configuration Audit Interaction

The Active UP performs external audit and PFD configuration audit interaction. The Session Manager gets a start and complete notification of the PFD configuration audit. The Session Manager does not start the external audit if the PFD configuration audit is in progress. If the PFD 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 the PFD audit is already underway.

Zero Accounting Loss for User Plane

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

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

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

Early PDU Recovery for UP Session Recovery

Early PDU Recovery feature overcomes the earlier limitation of Session Recovery feature wherein it did not 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 very long if there were a lot of 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

Prior to this release, the Session Recovery function did not 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 Prioritisation during Recovery, a separate skiplist 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 Control Plane and it is recovered first followed by normal calls.

These prioritized sessions also take priority in case of early PDU handling. The early PDU of normal calls will only initiate recovery when all prioritized sessions have been 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 Control Plane is responsible to set the priority flags for all the calls. The User Plane uses the priority call details received from the Control Plane for the Session Prioritisation feature.

Configuring 1:1 User Plane Redundancy for 4G CUPS

The following sections provide information about the CLI commands available to enable or disable the feature.

Configuring BFD Monitoring Between Active UP and Standby UP

Use the following commands to configure Bidirectional Forwarding Detection (BFD) monitoring on the Active UP and Standby UP. This command is configured 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 }
        exit
```

NOTES:

- **no**: Disables BFD monitoring on the Active and Standby UP.
- **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 Do not use the **chassis-to-router** keyword for BFD monitoring on the SRP link between the Active UP and the Standby UP.

- This command is disabled by default.

Flagging BGP Monitoring Failure

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

Note

- In this release, the **exclusive-failover** keyword is added to the existing **monitor bgp** CLI command as an alternate (new) algorithm to flag BGP monitoring failure.
- For more informnation about the **monitor bgp** CLI command in the "Service Redundancy Protocol Configuration Mode Commands" section command of the Command Reference Guide.
- Before adding the **exclusive-failover** keyword to the existing **monitor bgp** CLI command, implementing the **monitor bgp** command resulted in the following behavior:
  - BGP peer group was up if any BGP peer in that group was up.
  - Omitting a group configuration for a BGP monitor included that monitor in group 0.
  - BGP group 0 monitored in a context from an implicit group. Each context formed a separate BGP group 0 implicit monitor group.
  - BGP monitor was down if any BGP peer group was down.

```
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.
  - Removing a BGP peer being monitored induces a BGP monitor failure.
- This command is disabled by default.

Configuring Sx Monitoring on the Active UP and Standby UP

Use the following commands to configure Sx monitoring on the Active UP and Standby UP. 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 } } ]
         exit
```

NOTES:

- **no**: Disables Sx monitoring on the Active and Standby UP.
- **context***context_name* : Specifies the context of the Sx 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 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 peer, entered using IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.
- **disallow-switchover-on-peer-monitor-fail** :

  Prevents the switchback of the UP to Active state when the working status of the UP to CP link is unknown.
- It is possible to implement this CLI command multiple times for monitoring multiple Sx connections.
- The Sx monitor state goes down when any of the monitored Sx connections are down.
- This command is disabled by default.

Configuring SRP over IPSec on the Active UP and Standby UP

IPSec is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways. IPSec provides confidentiality, data integrity, access control, and data source authentication to IP datagrams.

The CUPS architecture uses the IPSec protocol to encrypt the packets sent over the Interchassis Session Recovery (ICSR) connection between the active and standby UPs. This encryption is done by defining an access-list to match all traffic between Service Redundancy Protocol (SRP) peers and associating it with a crypto map. This crypto map is used to establish Security Association between IPSec peers residing in UPs.

Note For more information on IPSec, its features or functionality, and applicable CLI configurations, refer the StarOS *IPSec Reference*.

The following CLI command is a sample configuration to configure SRP over IPSec for UPs.

```
context srp
   ip access-list srp-acl
   permit tcp host 209.165.200.225 host 209.165.200.226
   #exit
   ipsec transform-set A-foo
   #exit
   ikev2-ikesa transform-set ikesa-foo
   #exit
   crypto map srp-cm ikev2-ipv4
   match address srp-acl
   authentication local pre-shared-key key local key
   authentication remote pre-shared-key key remote key
   ikev2-ikesa transform-set list ikesa-foo
   payload foo-sa0 match ipv4
   ipsec transform-set list A-foo
   #exit
   peer 209.165.200.227
   #exit
   service-redundancy-protocol
   checkpoint session duration non-ims-session 30
   checkpoint session duration ims-session 30
   route-modifier threshold 18
   delta-route-modifier 2
   audit periodicity 60
   priority 2
   monitor bgp context isp 209.165.200.228
   monitor sx context EPC2 bind-address bbbb:abcd::77 peer-address bbbb:abcd::10
   peer-ip-address 209.165.200.226
   bind address 209.165.200.225
   #exit
   interface ike-lb loopback
   ip address 209.165.200.228 255.255.255.224
   crypto-map srp-cm
   #exit
   interface srp-rtr
   ip address 209.165.200.229 255.255.255.224
   #exit
   interface srp-loopback loopback
   ip address 209.165.200.225 255.255.255.224
   #exit
   ip route 209.165.200.226 255.255.255.224 209.165.200.231 srp-rtr
   ip route 209.165.200.227 255.255.255.224 209.165.200.231 srp-rtr
   #exit
```

Note IKEv1 - Transport mode with Authentication Header (AH) protocol is not recommended. Encapsulating Security Payload (ESP) is recommended because ESP performs both Authentication and Encryption.

Configuring VPP Monitor on Active UP and Standby UP

Use the following commands to configure Vector Packet Processing (VPP) monitor to trigger UP switchover on the Active UP 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 0-300 seconds
         exit
```

```
no monitor system vpp
```

NOTES:

- **no**: Disables VPP monitoring on the Active and Standby UP.
- **vpp delay-period***0-300 seconds* : Specifies the delay period in seconds for a switchover, after a VPP failure.

  If the delay period is a value greater than zero, 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 Switchback

Use the following commands to prevent the switchback of the new Standby UP to Active state again due to Sx monitor failure in the new Active. This command is configured in the SRP Configuration Mode.

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

Use either of the following CLIs to allow switchback of the new Standby UP 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 UP to Active state when the working status of the UP to CP link is unknown.

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

Note Assigning 0 seconds as the the timeout allows unplanned switchover.

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

- The default configuration is to allow unplanned switchover due to Sx 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 UP 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 Monitor Failure

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

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

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 4G CUPS 1:1 UP Redundancy feature:

- Type
- State
- 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 4G CUPS 1:1 UP Redundancy feature:

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

show srp monitor sx

The output of this CLI command contains the following fields for the 4G CUPS 1:1 UP Redundancy feature:

- Type
- State
- 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 4G CUPS 1:1 UP Redundancy feature:

- Type

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