GTPC and Sx Path Management

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	cnSGW-C
Applicable Platform(s)	SMI
Feature Default Setting	GTPC and Sx Path Management: Disabled – Configuration required to enable GTPC Path Failure: Enabled – Always-on Sx Path Failure: Enabled – Always-on Path Failure Detection Customization: Disabled – Configuration required to enable
Related Documentation	Not Applicable

Revision History

Table 2: Revision History

Revision Details	Release
Introduced support for IPv6.	2022.04.0
First introduced.	2021.02.0

Feature Description

The GTPC and Sx Path Management feature supports the following:

- GTPC path management using Echo Request and Echo Response messages.
- Sx path management using PFCP Heartbeat Request and Heartbeat Response. Node-level heartbeat procedures between the SGW-C and UPF.
- Detection of the GTPC path failure on S11 and S5 interface.
- Detection of the Sx path failure on the Sx interface.
- Configuration of the path failure detection policy to configure the path failure detection capability.

GTPC and Sx Path Management

Feature Description

GTPC and Sx Path Management supports the following:

- GTPC path management using Echo Request and Echo Response exchange over S5 and S11 interface to check peer aliveness.
- Sx path management using Packet Forwarding Control Protocol (PFCP) Heartbeat Request and Heartbeat Response exchange over Sx interface to check peer aliveness.

Feature Configuration

Configuring this feature involves the following steps:

- Configure the echo parameters. For more information, refer to Configuring the Echo Parameters, on page 3.
- Configure the heartbeat parameters. For more information, refer to Configuring Heartbeat, on page 3.
- Verify the peer configuration. For more information, refer to Viewing the Peer Configuration, on page 4.

Configuring the Echo Parameters

To configure the Echo parameters, use the following configuration:

Enabling the Echo Request

To enable the Echo Request, use the following configuration:

```
config
   instance instance-id instance_id
      endpoint endpoint_name
         interface [ s11 | s5e ]
           echo interval interval_value
           echo max-retransmissions max_retransmissions_count
           echo retransmission-timeout retransmission_timeout_count
            end
```

NOTES:

- **interval** *interval_value*—Specify the echo interval in seconds. Must be an integer in the range of 60-3600. Default value is 60 seconds.
- **max-retransmissions** *max_retransmissions_count*—Specify the maximum number of retries for GTP Echo Request. Must be an integer in the range of 0-15. Default value is 3.
- **retransmission-timeout** *retransmission_timeout_count*—Specify the Echo Request retransmission timeout period in seconds. Must be an integer in the range of 1-20. Default value is 5.

Disabling the Echo Request

To disable the Echo Request, use the following configuration:

```
config
   instance instance-id instance_id
      endpoint endpoint_name
         replicas replicas_count
         interface interface_name
            no echo
             end
```

Configuring Heartbeat

To configure the heartbeat parameters, use the following configuration:

Enabling Heartbeat

To enable a heartbeat, use the following configuration:

```
config
   instance instance-id instance_id
      endpoint pfcp
         interface sxa
            heartbeat
            interval interval
            retransmission-timeout timeout
```

```
            max-retransmissions retransmission_count
            end
```

NOTES:

- **interval** *heartbeat_interval*—Specify the heartbeat interval in seconds. Must be an integer in the range of 0-3600. To disable, set to 0.
- **max-retransmissions** *max_retransmissions*—Specify the maximum number of retries for the PFCP Heartbeat Request. Must be an integer in the range of 0-15. Default value is 4.
- **retransmission-timeout** *retransmission_timeout*—Specify the heartbeat retransmission timeout period in seconds. Must be an integer in the range of 1-20. Default value is 5.

Disabling Heartbeat

To disable a heartbeat, use the following configuration:

```
config
   instance instance-id instance_id
      endpoint pfcp
         interface sxa
            heartbeat
            interval interval
            end
```

NOTES:

- **interval** *heartbeat_interval*—Specify the heartbeat interval as 0 to disable the heartbeat.

Viewing the Peer Configuration

To view the peer restart counter, use the following configuration:

The following command displays the peer configuration:

```
show peers all [ endpoint ] [ local addr ] [ peer addr ]
```

```
show peers all SXA 209.165.201.12:8805 209.165.201.18:8805 POD CONNECTED
ENDPOINT LOCAL ADDRESS PEER ADDRESS DIRECTION INSTANCE TYPE TIME RPC ADDITIONAL DETAILS
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------
SXA 209.165.201.12:8805 209.165.201.18:8805 Inbound nodemgr-0 Udp 4 hours SGW-U Capacity:
65535,
LoadMetric: 0,LoadSeqNo: 0,Mode: Online,OverloadMetric: 0,OverloadSeqNo: 0,Priority: 65535

show peers all S11 209.165.201.4:2123 209.165.201.7:2123 LOCAL POD CONNECTED
ADDITIONAL ENDPOINT ADDRESS PEER ADDRESS DIRECTION INSTANCE TYPE TIME RPC DETAILS
-------------------------------------------------------------------------------------
-----------------
S11 209.165.201.4:2123 209.165.201.7:2123 Inbound nodemgr-0 Udp 25 seconds MME Recovery:
10

show peers all S5E 209.165.201.4:2123 209.165.201.21:2123 LOCAL POD CONNECTED
ADDITIONAL ENDPOINT ADDRESS PEER ADDRESS DIRECTION INSTANCE TYPE TIME RPC DETAILS
-------------------------------------------------------------------------------------
----------------
S5E 209.165.201.4:2123 209.165.201.21:2123 Inbound nodemgr-0 Udp 25 seconds PGW Recovery:
10
```

```
show peers all POD CONNECTED
ENDPOINT LOCAL ADDRESS PEER ADDRESS DIRECTION INSTANCE TYPE TIME RPC ADDITIONAL DETAILS
-------------------------------------------------------------------------------------------------------------------------------
<none> 209.165.201.29 209.165.201.18:8001 Outbound rest-ep-0 Rest 17 hours UDM <none>
<none> 209.165.201.29 209.165.201.18:8002 Outbound rest-ep-0 Rest 17 hours AMF <none>
<none> 209.165.201.29 209.165.201.18:8003 Outbound rest-ep-0 Rest 17 hours PCF <none>
<none> 209.165.201.29 209.165.201.18:8004 Outbound rest-ep-0 Rest 17 hours CHF <none>
<none> 209.165.201.29 209.165.201.18:9040 Outbound rest-ep-0 Rest 17 hours CHF <none>
S11 209.165.201.4:2123 209.165.201.6:2123 Inbound nodemgr-1 Udp 18 minutes MME Recovery:
10
S5E 209.165.201.12:2123 209.165.201.24:2123 Inbound nodemgr-1 Udp 5 hours PGW Recovery:
65535
SXA 209.165.201.12:8805 209.165.201.18:8805 Inbound nodemgr-0 Udp 22 minutes SGW-U Capacity:
 65535,LoadMetric: 0,LoadSeqNo: 0,Mode: Online,OverloadMetric: 0,OverloadSeqNo: 0,Priority:
 65535
```

Configuration Example

The following is an example configuration to enable the echo.

```
config
   instance instance-id 1
      endpoint gtp
         interface s11
            echo interval 60
            echo max-retransmissions 5
            echo retransmission-timeout 4
            end
```

The following is an example configuration to disable the echo.

```
config
   instance instance-id 1
      endpoint gtp
         replicas 1
         interface s5e
            no echo
            exit
         interface s11
            no echo
            end
```

OAM Support

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

Alerts

To configure Alerts for Peer Up and Peer Down, see *Key Performance Indicators* chapter in *Cisco Ultra Cloud Serving Gateway Control Plane Function - Metrics Reference*.

Bulk Statistics Support

Node Manager

The following are examples of Echo Transmitted and Echo Retransmitted messages:

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_retx",
gtpc_peer_ip="209.165.201.11",instance_id="1",interface_type="S5E",service_name="nodemgr"}
 3
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_tx",
gtpc_peer_ip="209.165.200.230",instance_id="1",interface_type="S11",service_name="nodemgr"}
 2
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_tx",
gtpc_peer_ip="209.165.201.11",instance_id="1",interface_type="S5E",service_name="nodemgr"}
 4
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_tx_initial",
gtpc_peer_ip="209.165.200.230",instance_id="1",interface_type="S11",service_name="nodemgr"}
 2
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_tx_initial",
gtpc_peer_ip="209.165.201.11",instance_id="1",interface_type="S5E",service_name="nodemgr"}
 1
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_res_rx",
gtpc_peer_ip="209.165.200.230",instance_id="1",interface_type="S11",service_name="nodemgr"}
 2
```

GTPC-EP Pod

The following are examples of Echo Request received and Echo Response sent messages:

```
gtpc_echo_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_rx",
gtpc_peer_ip="209.165.200.230",instance_id="0",service_name="gtpc-ep"} 1
```

```
gtpc_echo_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_rx",
gtpc_peer_ip="209.165.200.231",instance_id="0",service_name="gtpc-ep"} 1
```

```
gtpc_echo_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_req_rx",
gtpc_peer_ip="209.165.201.11",instance_id="0",service_name="gtpc-ep"} 1
```

```
gtpc_echo_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_res_tx",
gtpc_peer_ip="209.165.200.230",instance_id="0",service_name="gtpc-ep"} 1
```

```
gtpc_echo_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_res_tx",
gtpc_peer_ip="209.165.200.231",instance_id="0",service_name="gtpc-ep"} 1
```

```
gtpc_echo_msg_stats{app_name="smf",cluster="cn",data_center="cn",gtpc_msg_type="gtpc_echo_res_tx",
gtpc_peer_ip="209.165.201.11",instance_id="0",service_name="gtpc-ep"} 1
```

Procedure-Level

The following are examples of how to check the incremented values of Heartbeat Request, Heartbeat Response, and Heartbeat Request retry.

```
nodemgr_up_hb_msg_stats{app_name="smf",cluster="Local",current_nodemgr_id="0",data_center="DC",
instance_id="0",interface_type="SXA",primary_nodemgr_id="0",service_name="nodemgr",up_ep_key=
"209.165.201.1:209.165.201.21",up_msg_type="up_heartbeat_req_retx"} 3
```

```
nodemgr_up_hb_msg_stats{app_name="smf",cluster="Local",current_nodemgr_id="0",data_center="DC",
instance_id="0",interface_type="SXA",primary_nodemgr_id="0",service_name="nodemgr",up_ep_key=
"209.165.201.1:209.165.201.21",up_msg_type="up_heartbeat_req_tx"} 5
```

```
nodemgr_up_hb_msg_stats{app_name="smf",cluster="Local",current_nodemgr_id="0",data_center="DC",
instance_id="0",interface_type="SXA",primary_nodemgr_id="0",service_name="nodemgr",up_ep_key=
"209.165.201.1:209.165.201.21",up_msg_type="up_heartbeat_rsp_rx"} 5
```

Uniformity in compression at N4 and Sx interfaces

This feature introduces a peer-level capability exchange mechanism during the PFCP Association Setup procedure. It ensures that compression capabilities are correctly negotiated between the Control Plane (cnSGW-C) and the User Plane Function (UPF), preventing call rejections that occur in converged core deployments when a shared UPF has compression enabled.

Sx protocol compression is enabled by default in CUPS deployments. However, it is not supported on the converged core (cnSGW-C/SMF). When a UPF is shared between a converged core and a CUPS core, call rejections occur because the converged core cannot process compressed PFCP messages. This feature allows the cnSGW-C to explicitly declare its compression support status, ensuring the UPF sends uncompressed messages to converged core peers.

How it works

The feature functions through an exchange of Information Elements (IEs) during the PFCP Association Setup:

1. UPF Capability: The UPF sends the Operator configurable UPF capability IE (Type 359, as defined in 3GPP TS 29.244) in the Association Setup Request.
2. CP Capability: The cnSGW-C (CP) responds with a proprietary Operator configurable CP capability IE (Type 32769, Enterprise ID 9) in the Association Setup Response.
3. Capability Declaration: The cnSGW-C populates this IE with a value of 0 (Ccmpression disabled) to ensure the UPF sends uncompressed messages to this peer.

Note This feature is "Always-on" and auto-negotiated. No manual configuration is required to enable the capability exchange.

Per-Peer capability logic

The UPF determines whether to send compressed or uncompressed messages based on the declaration received from the SMF:

- IE 32769 (Value 0): Sent by SMF/cnSGW → UPF sends uncompressed messages.
- IE 32769 (Value 1): Sent by CUPS-CP → UPF sends compressed messages.

Monitor protocol

To verify the capability exchange, monitor the PFCP Association Setup Response sent by the SMF/cnSGW:

- IE Type: 32769
- Enterprise ID: 9
- Value: 0 (Compression disabled/not supported)

GTPC Path Failure

Feature Description

GTPC path failure detects peer-level GTPC path failure on the S11 and the S5 interface when:

- Echo Response contains a new restart counter value.
- Echo Request contains a new restart counter value.
- Echo Response is not received.
- Create Session Request or Modify Bearer Request contains a new restart counter value.
- Create Session Response or Modify Bearer Response contains a new restart counter value.

The connections may get disconnected due to different path failure is as follows:

- s11_path_failure
- s5e_path-failure
- s11_path_failure_local_purge
- s5e_ path_failure_local_purge
- s5e_recovery
- s11_recovery
- s5e_recovery_local_purge
- s11_recovery_local_purge

How it Works

This section describes how this feature works.

GTPC Path Failure Detection

Path failure is detected in the following conditions:

- Echo Failure: Echo failure occurs when the peer doesn't respond to the Echo Request or the retries.
- Restart Counter in Echo Response or Control Messages: The GTPC entity receives Recovery IE either in an Echo Response or from the peer GTPC message. GTPC entity compares the received the restart counter value with the previously stored restart counter value for that peer entity and performs the following:
  - Stores the received restart counter value for the peer when previously stored value isn't available.
  - When the max-remote-rc-change parameter is not configured, GTPC detects the change in the restart counter.
  - When max-remote-rc-change is configured, calculate the difference in the restart counter value considering restart counter rollover. Detects path failure when the difference between new and old restart counter is less than the value of max-remote-rc-change.

Note For more information on max-remote-rc-change, refer to Customization of Path Failure Detection, on page 13.

Path Failure Handling

Upon detecting a path failure, the network node notifies the failure through the Operation and Maintenance system and performs the following:

- Deletes the PDN connections (EPS bearer contexts) or the associated PDP contexts with peer IP address.
- Specifies the following actions for the selected interface:
  - **Local Purge:** The cnSGW-C clears the affected bearer (or PDN if the default bearer receives the path failure) locally without informing the peer. This action is default for all interfaces.

Note cnSGW-C sends the Sx Session Delete Request to UPF to clear session on path failure detection.

  - **Signal-Peer:** The cnSGW-C sends control signal towards the peer MME and P-GW.

    When signaling:
    - For PDN deletion, the SGW sends a Delete Session Request message to the PGW and a Delete Bearer Request (with LBI) message to the MME.
    - SGW sends a Delete Request on the S11 or the S5 interface to notify the peer.

Note Echo Request exchange is stopped when the peer is deleted.

Feature Configuration

Configuring this feature involves the following steps:

- Configure the action that must be taken on path failure detection. For more information, refer to Configuring Action on Path Failure Detection, on page 10.
- Configure the notification to update the peer node. For more information, refer to Configuring Notification to Update the Peer Node, on page 10.
- Verify the configuration. For more information, refer to Configuration Example, on page 10.

Configuring Action on Path Failure Detection

To configure the action for path failure detection, use the following configuration:

```
config
  profile sgw sgw_name
   path-failure [ s11 | s5e ] [ local-purge | signal-peer ]
   end
```

Configuring Notification to Update the Peer Node

Whenever cnSGW-C is restarted, the restart counter needs to be updated. For implementing this functionality, verify the Kubernetes use-volume-claims parameter value is set as true in Ops Center.

This configuration updates the restart counter when cnSGW-C restarts with the CLI system mode shutdown and system mode running.

Configuration Example

The following is an example configuration of path failure detection:

```
config
profile sgw sgw1
 path-failure s11 local-purge
 path-failure s5e local-purge
exit
```

OAM Support

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

Bulk Statistics Support

The following are examples of the GTPC path failure:

```
nodemgr_gtpc_pathfail_reasons{app_name="smf",cluster="cn",
data_center="cn",instance_id="1",pathfail_reason="pathfail_no_echo_rcv",
service_name="nodemgr"} 2
```

```
nodemgr_gtpc_pathfail_reasons{app_name="smf",cluster="cn",data_center="cn",
instance_id="0",pathfail_reason="pathfail_echo_res_rc_change",
service_name="nodemgr"} 1
```

Sx Path Failure

Feature Description

Sx path failure detects the path failure on the Sx interface when:

- Heartbeat Request contains a higher value of recovery timestamp.
- Heartbeat Response contains a higher value of recovery timestamp.
- Heartbeat Response is not received.
- Sx Association Request is received.
- cnSGW-C receives the Sx Association Update Request to release the peer.

How it Works

This section describes how this feature works.

Heartbeat Request

The cnSGW-C or UPF sends the Heartbeat Request on a path to the peer node to find out if the node is alive. The Heartbeat Request messages are sent for each peer with which a PFCP control association is established. cnSGW-C or UPF is prepared to receive the Heartbeat Request and it responds with a Heartbeat Response. The Heartbeat Request starts with the peer when a new session is established with the peer and it's stopped when the last session is released from the peer.

cnSGW-C and UPF send the Heartbeat Request based on the configured interval. If the peer doesn't respond, the message is retried for the configured number of times within the retry interval. After the response is received the defined action is taken for the calls associated with the corresponding peer.

Recovery Time Stamp is the IE which contains the start time of the peer node. The Heartbeat Request contains the selfrecovery timestamp value sent to the peer.

Note The heartbeat request is stopped only when the peer is deleted.

Heartbeat Response

The Heartbeat Response message is sent as a response to a received Heartbeat Request.

Recovery Timestamp is the IE which contains the start time of the node. Heartbeat Response contains the peer's Recovery Timestamp value.

Sx Path Failure Detection

Sx path failure is detected in the following conditions:

- **Heartbeat Failure:** When the peer doesn't respond to the heartbeat sent and also to the retries.
- **Recovery Timestamp Change in Heartbeat:** When the Heartbeat Response has a new Recovery Timestamp value then the previously received value. If the Recovery Timestamp value received is lower than the previously received value, the path failure isn't detected.
- **Sx Association Message:** When the Sx Association message is received again from the peer. In this case, all the calls are cleared and a notification is sent to eGTP peer.
- **Sx Association Release Message:** When the Sx Association release message is received. In this case, all the calls are cleared and a notification is sent to eGTP peer.

Path Failure Handling

When the recovery timestamp value received is more than the previously received value, the peer restart is detected. If the timestamp is lower than the previously received value, the value is ignored and peer restart isn't detected.

When the peer restart is detected to indicate the path failure for the peer, all the calls connected to that peer are cleared. The disconnection reason used for such calls is Sx path failure.

Sx association is also removed on detecting Sx path failure.

Heartbeat Handling

Whenever a PFCP entity receives a Heartbeat Request message (even from unknown peers), it responds with a Heartbeat Response message.

After a path failure is detected due to **No response to peer** error, no further Heartbeat Request is sent to that peer until the association is reestablished. Calls are cleared based on the path failure detection policy configuration.

Note After the Sx associations are removed, the heartbeat is stopped when Sx path failure is detected.

OAM Support

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

Bulk Statistics Support

The following are examples of the procedure-level statistics incremented for Heartbeat Request, Heartbeat Response, and Heartbeat Request retry:

```
nodemgr_up_hb_msg_stats{app_name="smf",cluster="Local",current_nodemgr_id="0",
data_center="DC",instance_id="0",interface_type="SXA",primary_nodemgr_id="0",service_name=
"nodemgr",up_ep_key="209.165.201.5:209.165.201.28",up_msg_type="up_heartbeat_req_retx"}
 3
```

```
nodemgr_up_hb_msg_stats{app_name="smf",cluster="Local",current_nodemgr_id="0",
data_center="DC",instance_id="0",interface_type="SXA",primary_nodemgr_id="0",service_name=
"nodemgr",up_ep_key="209.165.201.5:209.165.201.28",up_msg_type="up_heartbeat_req_tx"}
 5
```

```
nodemgr_up_hb_msg_stats{app_name="smf",cluster="Local",current_nodemgr_id="0",
data_center="DC",instance_id="0",interface_type="SXA",primary_nodemgr_id="0",service_name=
"nodemgr",up_ep_key="209.165.201.5:209.165.201.28",up_msg_type="up_heartbeat_rsp_rx"}
 5
```

Customization of Path Failure Detection

Feature Description

cnSGW-C lets you configure the path failure detection policy. By default, the path failure detection policy is enabled.

- **GTPC Path Failure Detection Customization:** GTPC path failure is detected when:
  - The Echo Request retries are exhausted.
  - The Echo Request or Response Restart counter is modified.
  - The control message Response Restart counter is modified.
  - If the absolute difference between the new and old restart counters is less than the value configured for max-remote-rc-change.

Note GTPC Path Failure Detection Customization allows user to ignore false peer restart with max remote restart counter (max-remote-rc-change) change functionality.

- **Sx Path Failure Detection Customization:** PFCP path failure is detected when:
  - The Heartbeat Request retries are exhausted.
  - The Heartbeat Request or Response recovery timestamps have modified.

Feature Configuration

Configuring this feature involves the following steps:

- Configure the GTPC path failure customization. For more information, refer to Configuring GTPC Path Failure Customization, on page 14.
- Configure the Sx path failure customization. For more information, refer to Configuring Sx Path Failure Customization, on page 14.

Configuring Sx Path Failure Customization

To configure the Sx path failure customization, use the following configuration:

```
config
  policy sx-path-failure-detection policy
   ignore heartbeat-retry-failure
   ignore heartbeat-recovery-timestamp-change
exit
 instance instance-id instance_id
   endpoint pfcp
     replicas replica_count
     sx-path-failure sx-detection-policy policy
     interface sxa
     sx-path-failure sx-detection-policy policy
     end
```

Configuring GTPC Path Failure Customization

To configure the GTPC path failure customization, use the following configuration:

```
config
   policy path-failure-detection policy_name
     max-remote-rc-change maximum_remote
     ignore echo-rc-change
     ignore control-rc-change
     ignore echo-failure
   exit
 exit
 instance instance-id instance_id
   endpoint gtp
     replicas replica_count
     vip-ip ipv4_address vip-port ipv4_port_number
     vip-ipv6 ipv6_address vip-ipv6-port ipv6_port_number
     dual-stack-transport { true | false }
     path-failure detection-policy policy
     interface [ s11 | s5e ]
     end
```

NOTES:

- When GTPC path failure detection policy isn't configured at interface-level, endpoint-level path failure detection policy is applicable.
- The max-remote-rc-change configuration specifies the counter change after which the S11 or S5 detects a peer restart. A peer restart is detected only if the absolute difference between the new and old restart counter is less than the value configured. For example, if the max-remote-rc-change is 10 and current peer restart counter is 251, then eGTP detects a peer restart only if the new restart counter is 252 through

255 or 0 through 5. Similarly, if the stored restart counter is 1, eGTP detects a peer restart only if the new restart counter is 2 through 11.

- Valid settings are from 1 to 255. The recommended setting is 32.

OAM Support

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

Bulk Statistics Support

GTPC Path Failure

Maintain statistics indicating number of times path failure was detected due to restart counter change in echo request or response message or control request or response message.

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="Local",data_center="DC",
gtpc_msg_type="gtpc_false_peer_restart_cfg_ctrl_rc_change",gtpc_peer_ip=
"209.165.201.17",instance_id="0",interface_type="S11",service_name="nodemgr"} 1
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="Local",data_center="DC",
gtpc_msg_type="gtpc_false_peer_restart_cfg_echo_rc_change",gtpc_peer_ip=
"209.165.201.17",instance_id="0",interface_type="S11",service_name="nodemgr"} 1
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="Local",data_center="DC",
gtpc_msg_type="gtpc_false_peer_restart_cfg_echo_rc_change",gtpc_peer_ip=
"209.165.201.27",instance_id="0",interface_type="S5E",service_name="nodemgr"} 1
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="Local",data_center="DC",
gtpc_msg_type="gtpc_ignore_echo_timeout",gtpc_peer_ip="209.165.201.27",
instance_id="0",interface_type="S5E",service_name="nodemgr"} 1
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="Local",data_center="DC",
gtpc_msg_type="gtpc_false_peer_restart_ignore_echo_rc_cfg",gtpc_peer_ip=
"209.165.201.27",instance_id="0",interface_type="S5E",service_name="nodemgr"} 1
```

```
nodemgr_gtpc_msg_stats{app_name="smf",cluster="Local",data_center="DC",
gtpc_msg_type="gtpc_false_peer_restart_ignore_ctrl_rc_cfg",gtpc_peer_ip=
"209.165.201.27",instance_id="0",interface_type="S5E",service_name="nodemgr"} 1
```

Table 3: GTPC Path Failure Statistics Descriptions

Statistics	Description
gtpc_false_peer_restart_cfg_echo_rc_change	The number of GTPC path failures ignored because Echo Restart Counter Change isn't within max-remote-rc-change configured.

Statistics	Description
gtpc_false_peer_restart_ignore_echo_rc_cfg	The number of GTPC path failures ignored because of Echo Restart Counter Change.
gtpc_false_peer_restart_cfg_ctrl_rc_change	The number of GTPC path failures ignored because Control Message Restart Counter Change isn't within max-remote-rc-change configured.
gtpc_false_peer_restart_ignore_ctrl_rc_cfg	The number of GTPC path failures ignored because of Control message Restart Counter Change.
gtpc_ignore_echo_timeout	The number of GTPC path failures ignored because of Echo Request timeout.

Sx Path Failure

Maintain statistics indicating number of times path failure was detected due to recovery timestamp change in the following messages.

```
nodemgr_up_pathfail_reasons{app_name="smf",cluster="cn",data_center="cn",
instance_id="0",service_name="nodemgr",up_pathfail_reason="up_pathfail_ignored_hb_retry"}
 1
```

```
nodemgr_up_pathfail_reasons{app_name="smf",cluster="cn",data_center="cn",
instance_id="1",service_name="nodemgr",up_pathfail_reason="up_pathfail_ignored_hb_rt_change"}
 1
```

```
nodemgr_up_pathfail_reasons{app_name="smf",cluster="cn",data_center="cn",
instance_id="1",service_name="nodemgr",up_pathfail_reason="up_pathfail_reason_association_release"}
 1
```

```
nodemgr_up_pathfail_reasons{app_name="smf",cluster="cn",data_center="cn",
instance_id="1",service_name="nodemgr",up_pathfail_reason="up_pathfail_reason_hb_retry"}
 8
```

```
nodemgr_up_pathfail_reasons{app_name="smf",cluster="cn",data_center="cn",
instance_id="1",service_name="nodemgr",up_pathfail_reason="up_pathfail_reason_hb_rt_change"}
 1
```

Table 4: Sx Path Failure Statistics Descriptions

Statistics	Description
up_pathfail_ignored_hb_retry	The number of Sx path failures ignored because of Heartbeat Request timeout.
up_pathfail_reason_hb_retry	The number of Sx path failures detected because of Heartbeat Request timeout.

Statistics	Description
up_pathfail_ignored_hb_rt_change	The number of Sx path failures ignored because of Heartbeat Request Recovery Timestamp Change Ignored.
up_pathfail_reason_hb_rt_change	The number of Sx path failures detected because of Heartbeat Request Recovery Timestamp Change.
up_pathfail_reason_association_release	The number of Sx path failures detected because of Sx Association Release.