

VLR Management

This chapter describes various MME features that provide additional resiliency of the Circuit Switched Fallback (CSFB) service, relating to the management of Visitor Location Registers (VLRs).

- Feature Summary and Revision History, on page 1
- Feature Description, on page 2
- Configuring VLR Offloading, on page 3
- Enabling UE Detach on VLR Failure or VLR Recover, on page 5
- Monitoring and Troubleshooting VLR Offload, on page 8

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	MME
Applicable Platform(s)	• ASR 5000 • ASR 5500 • VPC-DI • VPC-SI
Related Changes in This Release	Not Applicable
Related Documentation	• *Command Line Interface Reference* • *MME Administration Guide* • *Statistics and Counters Reference*

Revision History



Important Revision history details are not provided for features introduced before releases 21.2 and N5.1.

Revision Details	Release
Support is added for VLR Round-Robin load balancing.	21.27
First introduced	21.22

Feature Description

VLR Management is to manage the subscribers when a specific VLR is offloaded or failed. This chapter explains how MME distributes subscriber to available VLRs during VLR offload or failure.

These features require a valid license key to be installed. Contact your Cisco Account or Support Representative for information on how to obtain a license.

Passive VLR Offloading

The MME provides the ability for an operator to enable or disable "offload" mode for the specified Visitor Locations Registers (VLRs). This capability enables operators to preemptively move subscribers away from an SGs interface associated with a VLR which is planned for maintenance mode. When the **sgs offload**command is set on the MME, all sessions matching this VLR are marked with a "VLR offload" flag. During the next UE activity, the MME requires each UE to perform a combined TAU/LAU. This feature is available to all VLRs, both non-pooled VLRs and those configured within an MME LAC pool area.

The VLR offload functionality and MME offload functionality cannot be activated at the same time.

Active VLR Offloading

Active VLR Offloading provides all of the functionality of Passive VLR Offloading, but also actively detaches UEs associated with the VLR during an operator-specified time period. This expedites the process of offloading UEs prior to a planned VLR maintenance event. This feature is available to all VLRs, both non-pooled VLRs and those configured within an MME LAC pool area.

The VLR offload functionality and MME offload functionality cannot be activated at the same time.

UE Detach on VLR Recovery

The MME supports the ability to perform a controlled release of UEs when a failed VLR becomes active again. This feature is available to all VLRs, both non-pooled VLRs and those configured within an MME LAC pool area.

This applies to UEs that are currently registered as EPS-Only. This enables the UE to return to a combined attached state to restore SMS services.

UE Detach on VLR Failure

The MME supports the ability to perform a controlled release of UEs when an active VLR connection fails. This applies to CSFB UEs that are currently registered to the VLR that failed. This feature is available to all VLRs, both non-pooled VLRs and those configured within an MME LAC pool area.

This enables the UE to return to a combined attached state on a different VLR.

VLR Round-Robin Load Balancing

When multiple VLRs in the same pool are configured in the MME, and when one VLR fails, the subscribers on the failed VLR get attached to only one of the surviving VLRs. As a result, the subscribers are not distributed evenly across the surviving VLRs. The round-robin balancing algorithm is used to distribute subscribers across multiple available VLRs. VLR round-robin load balance is configurable through the CLI.

The MME allocates a VLR in the pool using the Round-robin algorithm for VLR offload (Active and Passive) and VLR failure cases. However, the MME does not use any form of load balancing if the default VLR is configured in pool area through **hash-value non-configured-value use-vlr** *vlr_name*. The default VLR setting takes precedence over the next available or round-robin load balance setting. By default, Round Robin Load Balancing is disabled.



Note The MME displays a warning message, if the default VLR is already configured while trying to configure the VLR Round-robin Load Balancing.

Configuring VLR Offloading

Enabling Passive VLR Offloading

The following Exec mode command instructs the MME to mark UEs associated with the specified VLR with a "VLR offload" flag. This enables the MME to preemptively move subscribers away from an VLR which is scheduled to be put in maintenance mode.

sgs offload sgs-service *service-name* **vlr** *vlr-name* **start time-duration 0 [-noconfirm]**

The following command stops the marking of subscribers associated with the specified VLR to an offload state.

sgs offload sgs-service *service-name* **vlr** *vlr-name* **stop [-noconfirm]**

Notes:

- A **time-duration** value of 0 enables Passive VLR Offloading only.
- More than one VLR may be offloaded at the same time.
- VLR Offloading and MME offloading cannot be performed at the same time.

Enabling Active VLR Offloading

The following Exec mode command instructs the MME to mark UEs associated with the specified VLR with a "VLR offload" flag, and begin detaching these UEs according to the time-duration specified in the command. Affected UEs are detached and required to reattach to another VLR.

sgs offload sgs-service *service-name* **vlr** *vlr-name* **start time-duration** *duration* **[-noconfirm]**

The following command stops active VLR offloading for UEs associated with the specified VLR.

sgs offload sgs-service *service-name* **vlr** *vlr-name* **stop [-noconfirm]**

Notes:

- A **start time-duration** *duration* entry must be an integer from 1 through 3000 to enables Active VLR Offloading and Passive VLR Offloading. The MME splits this time duration into *n* intervals, 5 seconds apart. A maximum of 50 subscribers will be actively detached per interval per session manager. For example, a setting of 5 minutes with 600 subscribers in a session manager (from the given VLR) would detach 10 subscribers per 5-second interval. Node level detach rate should be estimated by taking into account the number of session manager tasks. Any subscribers remaining at the expiry of the time-duration will not be detached, but will be marked with the "VLR offload" flag.
- VLR Offloading and MME offloading cannot be performed at the same time.

Configuring VLR Round-Robin Load Balance

Use the following configuration to enable VLR round-robin load balance:

```
configure
   context context_name
     sgs-service sgs_svc_name
          offload method { round-robin |  next-available }
         end
```

NOTES:

- **sgs-service** *sgs_svc_name*: Configures the SGs Service.
- **offload method { round-robin | next-available }**: Configures the round-robin or next-available load balancing method for affected sessions during VLR offload or failure.

Verifying VLR Offload Status and Configuration

The following command displays VLR offload statistics for the specified SGs service.

show sgs-service offload-status service-name *sgs_svc_name*

The following sample output shows VLR Offload related statistics.

```
show sgs-service offload-status service-name sgssvc
VLR Name                  : vlr1
VLR Offload               : Yes
Offloaded Count           : 31678
Total Count               : 43051
VLR Name                  : vlr2
VLR Offload               : No
```

```
Offloaded Count           : 0
Total Count               : 45789
```

To clear the counters displayed by the previous command, issue the following command.

clear sgs-service offload-status service-name *sgs_svc_name*

When Passive or Active VLR Offload is enabled, the following command displays the "VLR Offload" flag for the specified VLR.

show mme-service session vlr-name *vlr_name*

The following output shows the VLR Offload flag enabled.

```
show mme-service session vlr-name vlr1
   CSFB Information:
        SGS Assoc State:    SGs-ASSOCIATED
        SGS Service:        sgssvc
        VLR:                vlr1
        LAI:                123:456:200
        Pool Area:          pool1
        Non-Pool Area:      N/A
        P-TMSI:             0x1
        Flags:
   VLR Reliable Indicator
   VLR Offload
```

The following command shows the offload state of all VLRs on the system.

show sgs-service vlr-status full

```
show sgs-service vlr-status full
MMEMGR                        : Instance 6
MME Reset                     : Yes
Service ID                    : 2
Peer ID                       : 100794369
VLR Name                      : vlr1
SGS Service Name              : test
SGS Service Address           : 192.60.60.25
SGS Service Port              : 29118
VLR IP Address                : 192.60.60.6
VLR Psgsort                   : 29118
Assoc State                   : DOWN
Assoc State Up Count          : 2
VLR Offload                   : No
```

To clear the counters displayed by the previous command, issue either of the following commands. The first command clears statistics for all VLRs, while the second command clears statistics for the specified VLR only.

clear sgs-service vlr-status service-name *sgs_svc_name*
clear sgs-service vlr-status vlr-name *vlr_name*

Enabling UE Detach on VLR Failure or VLR Recover

UE Detach on VLR Recovery

The following Exec mode command instructs the MME to automatically perform active recovery of UEs when a failed VLR becomes responsive again.

sgs vlr-recover sgs-service *sgs_svc_name* **duration** *duration* **backoff-timer** *time* **[-noconfirm]**

Notes:

- When this command is issued, the MME monitors the availability of all VLRs. If a failed VLR become available again, the MME attempts to recover UEs that failed while the VLR was unavailable with an EPS Detach.
- When a VLR is down, and a UE needs to associate with the VLR that went down, the UE will be downgraded to EPS-Only-Attach when initially attaching. This command should be issued after the VLR recovers.
- UEs which required CSFB (voice) and were downgraded as a result of the VLR being down will not be affected by this command. This command remains active until it is disabled with the **no sgs vlr-recover** command.
- **duration** *duration* Specifies the number of minutes during which all qualifying UEs will be recovered.

  The MME splits this duration into *n* intervals, 5 seconds apart. A maximum of 50 subscribers are processed per interval per session manager. For example, a setting of 5 minutes with 600 subscribers in a session manager (from a given VLR) would result in the session manager processing 10 subscribers per 5-second interval. Node level detach rate should be estimated by taking into account the number of session manager tasks. Any subscribers remaining at the expiry of the duration will not be processed.
- **backoff-timer** *time* Specifies the number of seconds that the MME will wait, following the detection of a recovered VLR, before starting the VLR recovery actions.
- Refer to the *sgs vlr-recover* command in the Exec Mode chapter of the *Command Line Interface Reference* for more information.

The command listed below disables the **sgs vlr-recover** functionality.

no sgs vlr-recover sgs-service *sgs_svc_name* **[-noconfirm]**

UE Detach on VLR Failure

This functionality can be enabled manually, on an as-needed basis, using an Exec mode command, or it can be made a persistent configuration via an SGs Service Configuration Mode command. The following two sections describe how to configure each method (automatic and manual).

Important The MME will report a command line interface error (Invalid operation: VLR already set for failure.) if an attempt is made to configure/enable both methods simultaneously.

Configuring Automatic UE Detach on VLR Failure

The following commands configure the MME to automatically detect a VLR failure and initiate the controlled release of CSFB UEs. The configuration of this feature also allows a UE detach rate (UEs per second) to be defined.

```
configure
  context context_name
    sgs-service sgs_svc_name
```

```
        vlr-failure duration minutes backoff-timer seconds detach-rate number
[ -noconfirm ]
        end
```

The following commands disable this configuration:

```
configure
  context context_name
    sgs-service sgs_svc_name
      no vlr-failure [ -noconfirm ]
      end
```

Refer to the **vlr-failure** command in the *MME SGs Service Configuration Mode Commands* chapter of the *Command Line Interface Reference* for more information.

Manually Enabling UE Detach on VLR Failure

The following Exec mode command instructs the MME to perform controlled release of CSFB UEs connected to a VLR when a VLR becomes unavailable.

```
sgs vlr-failure sgs-service sgs_svc_name duration duratioin backoff-timer time
 [ -noconfirm ]
```

Notes:

- When enabled, the MME monitors the availability of all VLRs. If one or more VLRs become unavailable, the MME performs a controlled release (EPS IMSI detach) for all UEs associated with that VLR. If another VLR is available, the MME sends a combined TA/LA Update with IMSI attach.
- **duration** *duration* Specifies the number of minutes during which all qualifying UEs will be detached. Enter an integer from 1 to 3000.

  The MME splits this duration into *n* intervals, 5 seconds apart. A maximum of 50 subscribers are processed per interval by each session manager. For example, a setting of 5 minutes with 600 subscribers in a session manager (from a given VLR) would result in the session manager processing 10 subscribers per 5-second interval. Node level detach rate should be estimated by taking into account the number of session manager tasks. Any subscribers remaining at the expiry of the duration will not be processed.
- **backoff-timer** *time* Specifies the number of seconds the MME will wait following the detection of a VLR condition before starting the controlled release of affected UEs. Enter an integer from 1 through 3000.
- **detach-rate** This optional keyword specifies a maximum number of detaches to perform per 5 second cycle. **Note:** This keyword is available only for the **vlr-failure** command in the SGs Configuration Mode.

  For example, if 12,000 subscribers are to be detached during a 5 minute window (duration = 5 minutes), the MME calculates 60 cycles (5 minutes / 5-second cycles) which results in 200 UEs to detach per cycle.

  If the detach-rate is configured to 100, the MME will only detach 100 per 5 second cycle, resulting in a total of 6000 detaches. Any remaining UEs will remain attached until detached by other means (UE/network detach, etc).

The enabling command remains active until it is disabled with the following command:

```
no sgs vlr-failure sgs-service sgs_svc_name [ -noconfirm ]
```

Refer to the **sgs vlr-failure** command in the *Exec Mode (D-S)* chapter of the *Command Line Interface Reference* for more information.

Verifying UE Detach on VLR Failure/Recovery Status and Configuration

Use the following command to display the offload status of all VLRs on the system.

show sgs-service vlr-status full

This sample output shows the fields relating to UE Detach on VLR Failure and UE Detach on VLR Recover. Not all fields shown below may be displayed, based on your configuration:

```
show sgs-service vlr-status full
Exec Configured VLR Failure Detach : No         Detached Count : 0        Total : 0

SGs Service Configured
  VLR Failure Detach : Yes       Detached Count : 10     Total : 800
  VLR Recover Detach : Yes       Detached Count : 11     Total : 102
```

To clear the counters displayed by the previous command, issue either of the following commands. The first command clears statistics for all VLRs for the specified SG, while the second command clears statistics for the specified VLR only.

clear sgs-service vlr-status service-name *sgs_svc_name*
clear sgs-service vlr-status vlr-name *vlr_name*

Monitoring and Troubleshooting VLR Offload

SNMP Traps

The following traps are generated to track conditions relating to VLR associations:

The VLR down trap is raised only after the VLR goes to the DOWN state after being UP. When all VLR's are down after at least one has been UP, the all VLR's DOWN trap is raised.

- **starVLRAssocDown** and **starVLRAssocUp** - indicates a condition when an association of a VLR is down (VLRAssocDown), and when a down association comes back up (VLRAssocUp).
- **starVLRDown** and **starVLRUp** - indicates a condition where **all** SCTP associations to a specific VLR are down (VLRDown), and when a down VLR comes back up (VLRUp).
- **starVLRAllAssocDown** and **starVLRAllAssocDownClear** - indicates a condition when **all** SCTP associations of **all** VLRs are down (VLRAllAssocDown), and when a down association comes back up (VLRAllAssocDownClear).

Bulk Statistics

This SGs schema provides operational statistics that can be used for monitoring and troubleshooting the SGs connections on a per-VLR basis.

Refer to the *SGs Schema Statistics* chapter of the *Statistics and Counters Reference* for detailed explanations of all bulk statistics provided in this schema.

Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs.

VLR Offload Status

The following command shows the status of the VLR offload process for the specified SGs service.

show sgs-service offload-status service-name *sgs_svc_name*

The following command shows the status and configuration information of all VLRs on the system.

show sgs-service vlr-status full

UE Detach on VLR Recovery and VLR Failure

The following command shows the statistics resulting from the **sgs vlr-recover** and **sgs vlr-failure** commands.

show sgs-service vlr-status full

Refer to the *show sgs-service* chapter of the *Statistics and Counters Reference* for detailed explanations of all information displayed by this command.