

vMRF Overview

Virtual Multimedia Resource Function

Function Survey

Copyright

© Ericsson AB 2016–2019. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document *Trademark Information*.



Contents

1	Function Overview	1
1.1	vMRF Architecture	2
1.2	Virtual Machine Functions in vMRF	2
1.3	Roaming SC in vMRF	3
2	Audio Codec Selection Principles	5
3	Audio Announcements	6
4	Detection of Dual Tone Multi Frequency (DTMF) Tones	7
5	Forwarding of DTMF Tones through Audio Conference	8
6	Playing Tones	9
7	Emergency and Priority Call Handling	10
8	Audio Conferencing	11
9	Overload Protection	12
10	Shared Storage	13
11	Automatic Backup and Restore	14
12	Address Resolution and Next Hop Supervision	15
12.1	VLAN Transparency	16
13	Media Stream Processing Services	17
13.1	Adaptive Multi-Rate (AMR) Speech Coder Service	17
13.2	Adaptive Multi-Rate Wideband (AMR-WB) Speech Coder Service	17
13.3	Announcement Service	18
13.4	Audio Mixing Service	18
13.5	DTMF Receiver Service (DTMF-R)	18
13.6	DTMF Sender Service (DTMF-S)	18
13.7	Enhanced Voice Services (EVS)	19
13.8	G.729 Service (G.729)	19
13.9	Jitter Handling Service	19
13.10	Media Stream Recording	20



13.11	Pulse Code Modulation Service (PCM)	20
13.12	Real-Time Transport Protocol and Real-Time Transport Control Protocol Service (RTP/RTCP)	20
13.13	Tone Sender Service	20
13.14	User Plane Frame Handler Service (UP FH)	21
14	Support for Watchdog	22
15	VMware Tools	23
16	Limitations and Differences	24
16.1	vMRF VNF Size Limitations	24
	Reference List	25



1 Function Overview

The Virtual Multimedia Resource Function (vMRF) is used to mix and process media streams. vMRF is controlled by the Multimedia Telephony Application Server (MTAS), using H.248 control protocol. [Figure 1](#) illustrates signaling and payload traffic related to vMRF.

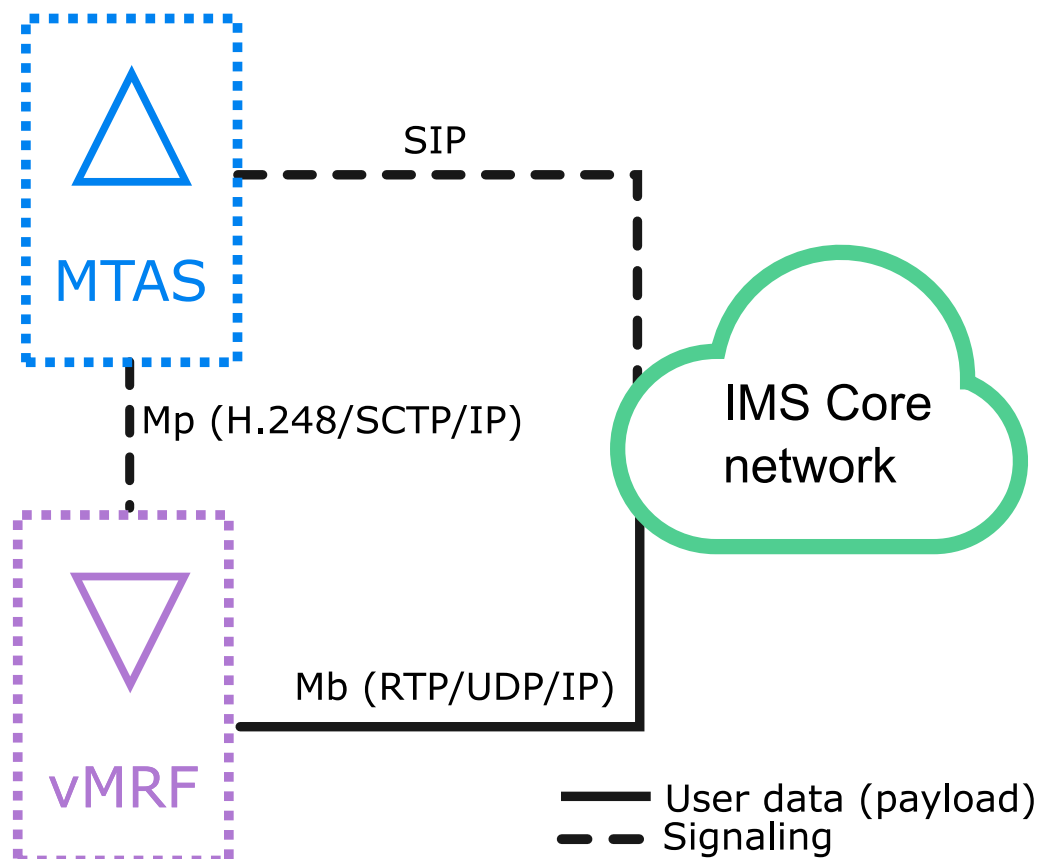


Figure 1 vMRF Overview

vMRF supports the following functions:

- Playing audio announcements
- Audio conferencing
- Detection of DTMF tones
- DTMF tone forwarding for audio conference
- Playing tones

vMRF supports the following audio codecs: AMR-NB, AMR-WB, EVS, G.711, G.722, and G.729.

Note: The Consumer Communication capacity license is needed for basic vMRF functions.

The use of EVS requires the Enhanced Voice Services capacity license and connection to a Network License Server (NeLS).

For more information on licensing, see [License Management](#).

1.1 vMRF Architecture

vMRF is a Virtualized Network Function (VNF). A single VNF contains multiple VMs. See [Figure 2](#) for an example overview of the vMRF.

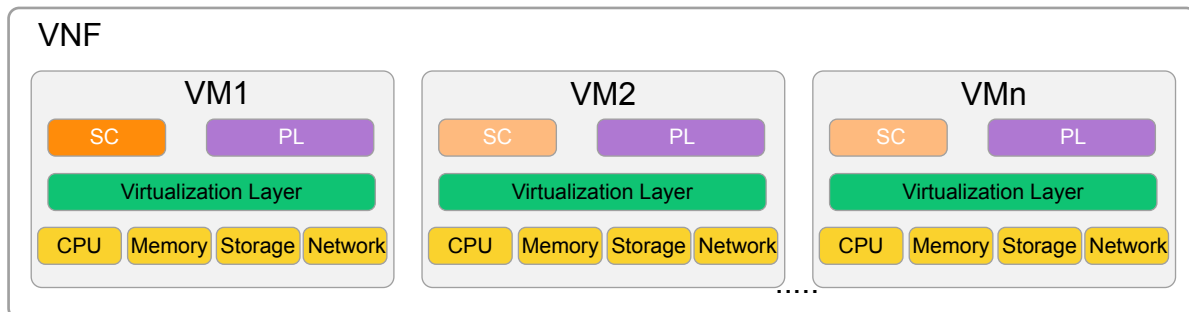


Figure 2 vMRF Architecture

The vMRF VNF can be deployed in the network multiple times, each time as a separate VNF.

1.2 Virtual Machine Functions in vMRF

In the vMRF VNF, each VM provides the following functions:

- Payload (PL) function: The PL function processes user plane traffic and H.248 signaling traffic. Each VM in a VNF provides the same PL functionality and has the same configuration (except local IP addresses).
- System Controller (SC) function: The SC function processes O&M traffic and is responsible for the VNF internal clustering function that is needed for scaling out and scaling in. All VMs in the VNF can act as the SC, but the SC function is active in only one VM at a time. The active SC VM is selected by the Roaming SC function, see [Roaming SC in vMRF](#) on page 3.

See [Figure 3](#) for an overview of the vMRF architecture.

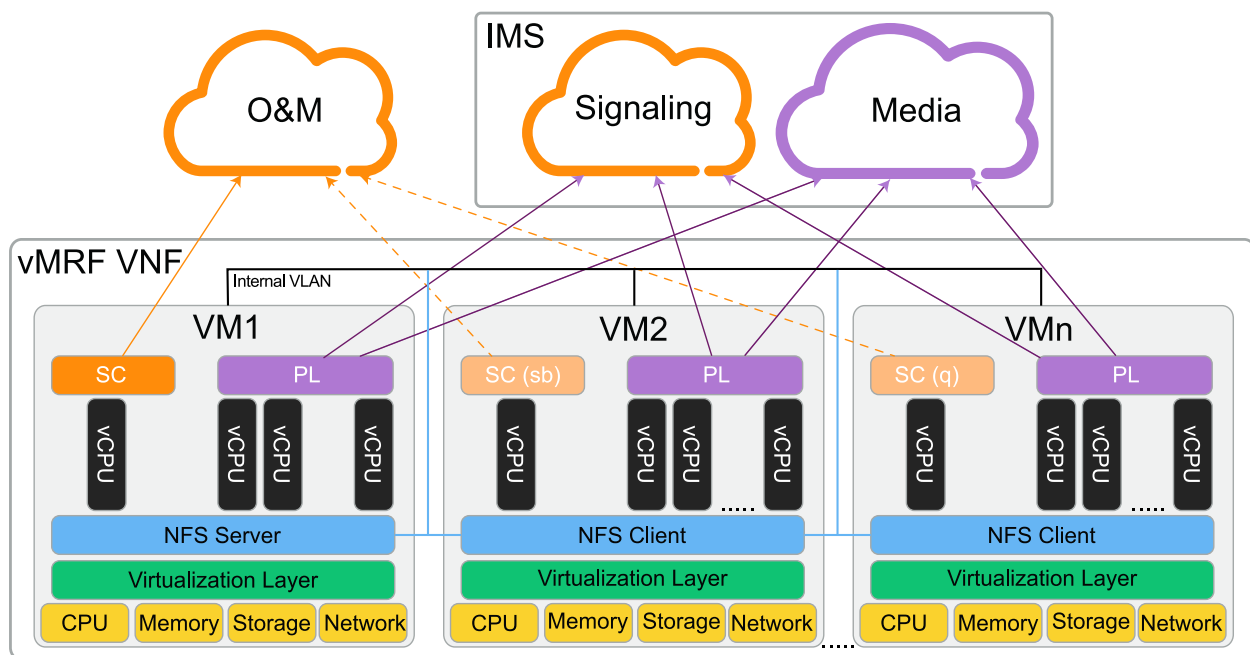


Figure 3 vMRF Architecture

Configuration data is synchronized between the VMs using internal NFS servers and clients connected through the internal VLAN. This VLAN must not be exposed to external networks. The active SC VM contains the NFS server, while the PL VMs the clients. Both server and clients connect to the shared location at `/media/cluster_sync`.

1.3 Roaming SC in vMRF

In vMRF, O&M functions are secured with a Redundancy scheme called Roaming SC. This means that all virtual machines can process payload and also act as a system controller when necessary. If the SC fails, any virtual machine in the cluster can take over the SC role. This provides the robustness required by operating the VNF in a cloud environment. This also allows the VNF to operate autonomously, eliminating the need for manual clustering-related configuration.

With the Roaming SC scheme, all SCs have one of the following states:

- **ACTIVE:** The ACTIVE SC is processing O&M traffic for the VNF. The SC function runs on a separate core from the PL functions.
- **STANDBY (sb):** The STANDBY SC functions only as a PL VM, but it becomes the ACTIVE SC in case of the failure of the current ACTIVE SC.
- **QUIESCED (q):** All VMs that are not in STANDBY or ACTIVE SC state are in QUIESCED state. They function as PL VMs, and are ready to take either the ACTIVE or STANDBY SC state when needed.

The SC state is indicated in the `scState` attribute of the `MrfInstance` MO as follows:



SC_ACTIVE	This VM is the active SC.
SC_STANDBY	This VM is the standby SC.
SC_QUIESCED	This VM is in QUIESCED state.

Assignment of the different SC states is done according to the following:

- **ACTIVE:** After a cluster restart or the simultaneous failure of the ACTIVE and STANDBY SC VMs, the ACTIVE SC is selected with a leader election algorithm.
- **STANDBY:** When the STANDBY SC fails or becomes ACTIVE, a new STANDBY SC is selected randomly from the QUIESCED SC VMs in the VNF.
- **QUIESCED:** All SCs that are not ACTIVE or STANDBY are in QUIESCED state.

When a VM takes the ACTIVE SC role, it automatically inherits the O&M IP address of the former ACTIVE SC. This ensures a seamless and transparent switchover of the O&M end-point.

Additionally, the new ACTIVE SC takes over the internal NFS server function. The NFS clients in the rest of the VMs mount this server.

To prevent two SCs operating at the same time, the new SC VM starts processing O&M traffic only after a delay period: about 5 seconds after cluster startup, or 6.5 seconds after the failure of the ACTIVE SC VM.



2 Audio Codec Selection Principles

For the incoming SDP offer the highest priority audio codec received in the SDP offer that is supported by the vMRF is selected for the media stream. In addition, if the SDP offer included also a telephone-event corresponding to the RTP clock rate of the selected audio codec, then the telephone-event is also selected for the media stream. The selected audio codec, and possibly the selected telephone-event, is sent back in the SDP answer.

For outgoing SDP offer the following audio codecs are sent in the SDP offer in this priority order: EVS, AMR-WB, G.722, AMR, G.729, PCMA, and PCMU. The highest priority audio codec, and possibly the telephone-event corresponding to the RTP clock rate of that audio codec, received in the SDP answer is used for the media stream.



3 Audio Announcements

Audio announcements feature is used to provide audio announcements and to facilitate interactive services for operators and network providers. vMRF supports audio announcements through streaming of audio content from the announcement files to subscribers. Audio announcements are supported in HD quality. The sample rate of the audio announcements is 16 kHz and the audio mode is Mono. Audio announcements have to be encoded in 16 bit 16 kHz PCM and stored in Waveform Audio File Format (WAV). vMRF plays an announcement towards a subscriber when a play request is received from the controlling node over H.248. Announcement files are stored on a storage server, and referenced by an announcement ID or variable announcement type. The recently played announcement files are cached to the local disk of a vMRF VM to support faster playing of announcements for future playing requests.

vMRF supports fixed announcements and segmented announcements. A fixed announcement consists of a single basic announcement. A segmented announcement is a compilation of individual announcements referred as announcement segments. In vMRF, an announcement segment can be a basic announcement or a stand-alone variable announcement. A basic announcement is a standalone announcement file that contains audio data of a spoken word or phrase. A variable announcement is an announcement whose spoken information heard by the user depends on input data received in a play request. vMRF supports the following variable announcement types: digits, date, time, number, and money. The variable announcement has a variable announcement logic file associated to each variable announcement type.

vMRF supports announcements with multiple languages. Multi-language announcements present the operator with an option to play announcements toward end users in their preferred language. Multi-language announcements are supported for both fixed announcements and segmented announcements. Multi-language announcements are identified by a single announcement ID or variable announcement type, and different languages are identified by using language tags. This enables playing of announcements toward end users in their preferred language, and simplifies the announcement configuration in vMRF and in the controlling node.

For configuration information on audio announcements, see [vMRF Configuration Management](#).



4 Detection of Dual Tone Multi Frequency (DTMF) Tones

DTMF digits can be detected either as inband signals from an audio stream or as separate telephone events. In case of a telephone event, the DTMF digits are detected from the RTP packets that use *telephone-event* payload format with dynamically established RTP payload type number. DTMF detection is done either for inband signals or for telephone event, not for both types simultaneously. If telephone event payload type has been negotiated for the audio stream, then vMRF detects DTMF telephone events, otherwise audio inband signals.

DTMF digits are specified in a dialing plan, called a digit map. The digit map resides in the vMRF and it is used to detect and report certain DTMF patterns received from the user plane. A pattern is a set of rules which express how the user is expected to dial DTMF digits, including, for example, digit names and timer values for the detection process. Up to 16 different DTMF digits can be used, corresponding to the following symbols: 0, 1, 2...9, A, B, C, D, *, #. The digit map is preloaded dynamically for each session by the controlling node over H.248. vMRF detects as many DTMF digits as it is contained in the digit map. If all digits match, vMRF notifies the controlling node that DTMF digits were collected successfully. The notification contains the digits that were collected and number of attempts that were made to detect the digits.



5 Forwarding of DTMF Tones through Audio Conference

vMRF supports forwarding of DTMF digits through audio conference from one participant to all other participants.

The function supports detection of either telephone-event DTMF or audio DTMF from the participant, based on SDP negotiation. While detecting audio DTMF digits, the received audio DTMF digits are filtered out from the stream before mixing audio by the audio mixer. The detected DTMF digits are sent to all other participants either as telephone-event DTMF or audio DTMF, based on SDP negotiation with each participant. The context topology and termination stream mode information is checked and obeyed before forwarding DTMF digits from one participant to another participant. The function can be enabled by configuration.



6 Playing Tones

vMRF supports playing of tones to the user plane based on request from the MTAS. For supported tones, see [Tone Sender Service](#) on page 20.



7

Emergency and Priority Call Handling

vMRF supports emergency and priority call handling for the following call types:

- Emergency calls
- Priority calls
- International Emergency Preference Scheme (IEPS) priority calls

A call is considered an emergency or priority call if it is indicated so in the call setup. Emergency and priority calls are treated the same way, and they are prioritized over non-emergency and non-priority calls in the case of high load on the vMRF.

A certain amount of processing capacity and internal resources are reserved by the vMRF for emergency and priority calls. This reserved fraction is called a priority pool. Normal calls are not allowed to use the priority pool. Emergency and priority calls are allowed to use resources outside of the priority pool.

The size of the priority pool can be configured in the vMRF. For more information on priority pool, see [vMRF Configuration Management](#).

For emergency and priority call handling the following rules apply:

- Emergency and priority calls are never rejected because of network level admission control or high CPU load
- Emergency and priority calls can be rejected because of lack of internal resources



8 Audio Conferencing

The audio conferencing feature provides a service that allows the connection of as many as 30 participants to a conference call. For each participant, it can be specified whether the connection is unidirectional (listening only) or bidirectional. The audio conferencing feature can dynamically mix audio from several participants. The audio from a participant is mixed to the audio conference based on active speaker logic.

Each client can have its own codec in use, and the Audio Conferencing feature provides transcoding between them. Audio mixing is performed at 16 kHz sampling rate. This guarantees wideband conferencing experience for wideband-capable terminals.

The following audio codecs are supported by vMRF: PCM, EVS, G.722, G.729, AMR-NB, and AMR-WB.



9 Overload Protection

To detect an overload situation in the network caused by high network traffic, the vMRF monitors the following:

- vMRF Application processor load (H.248 signaling, control of media stream processing)
- vMRF Media Stream Processing load (computationally demanding media processing, for example, transcoding)
- vMRF IP Pipeline load (simple media processing, for example, IP packet Header processing)
- vMRF IP Pipeline vSwitch load (outgoing packets are dropped if the vSwitch cannot process more packets)

In the case of an overload situation, an alarm is raised, based on available processing capacity. When 100% of the available processing capacity is used in the vMRF functions, the `MRF Instance Overloaded` alarm is raised. In this case only emergency and priority calls are accepted. VNF processing capacity must be increased, for example, by adding new VMs to avoid traffic loss. When a call is rejected because of overload, the H.248 ErrorCode 510 "Insufficient Resources" is reported.



10 Shared Storage

The use of shared storage is optional in vMRF. It allows for the storage of log files, crash dump files, and configuration backups on a remote server. To use it, a remote server with SSHFS (Secure Shell Filesystem) installed must be configured. The remote server and the exported files are managed by the operator.

The following files are stored by each VM:

- log files from the `/var/log` directory, including journal log files
- crash dump files
- configuration backup files created by the automatic backup and restore function, described in [Automatic Backup and Restore](#) on page 14

vMRF connects to the server with SSHFS (Secure Shell Filesystem), mounts a specified directory path, and creates a subfolder for the cluster, and subfolders for the files for each VM in the cluster. This ensures that logs and other shared files of different VMs and VNF instances do not get mixed up. For authentication, an SSH key pair has to be created. This key pair can be cluster-specific, or common for all clusters. The prerequisites for the remote server are the following:

- ssh connection support with public and private keys
- authentication of the ssh connections from the vMRF virtual machines
- defined username for the SSHFS mount
- ssh public key for the SSHFS mount
- the public key is appended to the authorized keys file using the following command:

```
cat .ssh/<public_key_file_name> >> .ssh/<authorized_keys_file_name>
```

- path for the files

Configuration of the shared storage in the vMRF is performed during deployment.

For more information on the HOT template configuration, see [Deployment Guide for OpenStack](#) and [Deployment Guide for Cloud Execution Environment \(CEE\)](#).

For OVF configuration, see [Deployment Guide for VMware vSphere](#) and [Deployment Guide for VMware vCloud Director](#).



11 Automatic Backup and Restore

vMRF performs automatic configuration backup after configuration changes. To prevent unnecessary backup operations, a new configuration is saved after an 80-second wait period following the most recent change. The 10 latest configuration backups are stored in time-stamped files in `/cluster/storage/configurationbackups/`. The file naming convention is the following:

```
<YYYYMMDD>_<hhmmss>_mrf.tar.gz
```

The copy of the latest backup file is also saved as `mrvs_config.tar.gz` in the same directory. This file is synchronized across all VMs of a cluster to enable any new SC to restore configuration after cluster restart.

When the maximum number of backup files is reached, the oldest backup file is deleted before a new backup is stored.

vMRF can perform automatic configuration restore if all VMs in a cluster reboot simultaneously. In this case, the new SC VM looks for `mrvs_config.tar.gz` in `/cluster/storage/configurationbackups/` and imports the configuration to the VNF, if the file exists.



12 Address Resolution and Next Hop Supervision

IPv4 addresses are 32 bits long. In written format, the address is divided into four 8-bit (1 byte) long sections, using dots as separators. For example, 172.16.254.1 (decimal digits).

IPv6 addresses are 128 bits long. In written format, the address is divided into eight 16-bit long sections (pairs of bytes), using colons as separators. For example, 2001:db8:ffff:1:201:2ff:fe03:405 (hexadecimal digits).

For alternative writing formats of IPv6 addresses, see [RFC 4291](#).

Local IP address allocation method depends on the used virtualization environment and user choice.. For more information on the available methods, see [vMRF Networks and Connectivity](#) and the relevant [Deployment Guides](#).

IPv4 and IPv6 use the following different methods for address resolution and next hop supervision:

- Address Resolution Protocol (ARP): It maps an IPv4 address to the corresponding MAC address.
- Neighbor Discovery (ND): It maps an IPv6 address to the corresponding MAC address.

The following subfunctions are supported by both the ARP and the ND:

- Next hop MAC address fetching
- Next hop (peer) supervision
- Address collision detection

Next hop addresses for the following networks are fetched from the HOT or OVF files after instantiation:

- IPv4 media networks
- O&M networks
- Signaling networks

Internet Control Message Protocol for IPv6 (ICMPv6) is used for fetching the next hop addresses.



12.1 VLAN Transparency

vMRF provides VLAN transparency by using VLAN tags for packets towards external networks in deployments using MO-based IP allocation with IP pools.

Note: VLAN transparency is not supported in VMware deployments.

If VLAN transparency is enabled, the OpenStack port does not overwrite the VLAN ID (configured in vMRF) in the sent packets.

VLAN IDs can be configured by the following attributes:

- Media interfaces: `mediaVlanId` in `MrfConfiguration`

Media VLAN ID has to be configured before the configuration of IP pools. The media VLAN ID cannot be changed unless all IP pools are removed.

- Signaling interfaces: `signalingVlanId` in `MrfH248Control`

Before setting the signaling VLAN ID, all `MrfH248Interface` MOs must be locked. Alternatively, the signaling VLAN ID can be set before creating the `MrfH248Interface` MOs.

These attributes are applicable only if the VLAN transparency feature has been enabled at deployment.

Traffic sent through the following interfaces remain untagged:

- Internal
- Internal O&M
- External O&M



13 Media Stream Processing Services

13.1 Adaptive Multi-Rate (AMR) Speech Coder Service

The Adaptive Multi-Rate (AMR) Speech Coder is a high-quality narrowband codec standardized by 3GPP. AMR is standardized as the mandatory speech coder for the LTE radio network; it is also used in WCDMA and GSM networks, and in Voice over IP solutions. The speech coder handles the conversion between PCM coded speech and AMR coded speech. Discontinuous Transmission (DTX) is supported in both uplink and downlink. DTX allows the radio transmitter to be switched off most of the time during speech pauses. This saves power in the UE and reduces interference over the air interface.

The AMR speech coder operates in different modes, representing source rates from 4.75 kbps to 12.2 kbps. The selected rate is a trade-off between speech quality and robustness. AMR complies with [3GPP TS 26.071](#).

13.2 Adaptive Multi-Rate Wideband (AMR-WB) Speech Coder Service

The Adaptive Multi-Rate Wideband (AMR-WB) codec is used to provide better speech quality than the existing Adaptive Multi-Rate Narrowband codec. AMR-WB provides better speech quality because its analog spectrum range is 100–7000 Hz instead of the 300–3400 Hz of narrowband codecs. The sampling rate for AMR-WB is 16 kHz instead of the 8 kHz of narrowband codecs. The prerequisite for better speech quality is that the speech is carried end-to-end in AMR-WB coded form between the two UEs. HD speech quality is lost if AMR-WB speech is converted to a narrowband codec. This is the case, for example, in calls between an AMR-WB mobile and a PSTN phone.

The AMR-WB speech codec consists of the following components:

- Multi-rate speech codec
- Source-controlled rate scheme, including voice activity detector
- Comfort noise generation system
- Error concealment mechanism (to combat the effects of transmission errors and lost packets)

The AMR-WB codec has the following nine source rates: 6.60, 8.85, 12.65, 14.25, 15.85, 18.25, 19.85, 23.05, and 23.85. These AMR-WB rates correspond to SDP mode-set values from 0 to 8 (no mode-set on SDP means the same as all rates



supported). The codec is able to switch its bit rate every 20 ms. See [3GPP TS 26.171](#).

13.3 Announcement Service

The Announcement Service plays an announcement towards a subscriber. The announcement is read from the announcement cache, where the announcement was fetched from the announcement storage when a play request was received. The Announcement Service supports playing of announcements that are encoded in 16 bit 16 kHz PCM and stored in Waveform Audio File Format (WAV).

13.4 Audio Mixing Service

The Audio Mixing service is responsible for audio mixing during audio conference.

The Audio Mixing service has an active speaker logic that determines which clients are included in the audio mix. This is based on voice activity detection. The detection algorithms calculate frame power, detect pitch, and estimate tonal activity to determine whether a client is speaking or not. Active speaker logic also ensures that the client has a long enough speech period before it is included in the audio mix, and a hangover period before the included client is dropped from the audio mix.

Audio mixing is performed at 16 kHz sampling rate.

13.5 DTMF Receiver Service (DTMF-R)

The DTMF-R detects DTMF digits from the user plane. Both Audio DTMF and Telephone-Event DTMF are supported. The DTMF-R can detect up to 16 different DTMF digits, corresponding to 0–9, A, B, C, D, *, and #.

For Audio DTMF, the DTMF-R detects each signaling tone, validating a correct tone pair and checking the timing. In the presence of speech, additional tests are performed to avoid interpreting a speech waveform as a valid signal tone. DTMF-R can be configured to detect overdriven DTMF tones. Without this setting, these tones would be dropped, because the harmonics produced by the low-frequency DTMF tones that fall within the passband of the high-frequency DTMF tones can be considered an indication of speech. For Telephone-Event DTMF, the digits are interpreted from RTP packets according to [RFC 4733](#).

13.6 DTMF Sender Service (DTMF-S)

The DTMF Sender (DTMF-S) sends DTMF digits to the user plane.

Both Audio DTMF and Telephone-Event DTMF are supported.



The DTMF-S can generate up to 16 different DTMF signals, corresponding to 0–9, A, B, C, D, *, and #. Audio DTMF consists of two simultaneously played frequencies. Telephone-Event DTMF is encoded to RTP according to [RFC 4733](#).

13.7 Enhanced Voice Services (EVS)

Enhanced Voice Services (EVS) is a multi-rate audio codec that operates at 8 kHz, 16 kHz, 32 kHz, and 48 kHz sampling rates, and offers full audio bandwidth ranging from 20 Hz up to 20 kHz. EVS supports bit rates from 5.9 kbps to 128 kbps. EVS supports comfort noise generation and error concealment.

The media can be encoded with EVS or with EVS AMR-WB IO mode. The frame size is 20 ms and the RTP timestamp is always generated with 16kHz clock. 20 ms and 40 ms packetization times are supported for EVS.

Note: The use of the EVS codec requires the Enhanced Voice Services capacity license and connection to a Network License Server (NeLS).

13.8 G.729 Service (G.729)

The G.729 high compression codec enables packet voice transmission using G.729 towards external IP-based networks. vMRF supports the G.729 speech codec according to [ITU-T G.729](#), with the following variants:

- G.729A, reduced complexity 8 kbps Conjugate Structure Algebraic Code Excited Linear Prediction (CS-ACELP)
- G.729AB, meaning G.729A with built-in silence suppression

vMRF supports G.729A with packet loss concealment. A packet loss concealment mechanism is supported in order to minimize the distortion in output speech caused by lost or excessively late speech packets. The G.729A is bit stream interoperable with the full version of G.729.

vMRF supports G.729AB, which adds on silence suppression including voice activity detection, generation and reception of SID (Silence Insertion Descriptor) packets, as well as comfort noise generation. These algorithms are used to reduce the transmission rate during silence periods of speech. The G.729AB (defining reduced complexity version of the G.729 speech codec with built-in silence suppression) is bit stream interoperable with the full version of G.729 with silence suppression.

13.9 Jitter Handling Service

Jitter (delay variation between data packets) is removed or decreased using jitter compensation performed by the Jitter Handling service. The Jitter Handling service supports adaptive jitter service.



In the beginning of the call the jitter buffer size is always the configured initial jitter buffer size, but during the call the jitter buffer size adapts to the measured jitter.

13.10 Media Stream Recording

The Media Stream Recording (MSR) functionality provides an efficient method for Ericsson support personnel to perform advanced troubleshooting on media plane related problems that are officially reported by the customer. MSR is used only after the customer has explicitly requested Ericsson to analyze the problem and has agreed with Ericsson on the use of MSR. Recordings can only be made by Ericsson personnel, and the produced recording files are sent to Ericsson Virtual Multimedia Resource Function (vMRF) support for analysis.

13.11 Pulse Code Modulation Service (PCM)

The Pulse Code Modulation (PCM) service handles coding and decoding of PCM ([ITU-T G.711](#)) samples in vMRF. Both A-law and μ -law formats are supported. PLC is supported for the PCM service.

13.12 Real-Time Transport Protocol and Real-Time Transport Control Protocol Service (RTP/RTCP)

The RTP/RTCP service represents the RTP and RTCP protocol termination for IP user plane traffic.

RTCP can be used to monitor the aliveness of the session as well as quality of service during an ongoing RTP session, according to [RFC 3550](#).

13.13 Tone Sender Service

The Tone Sender service supports the generation of the following tones:

- H.248.1 E.7 Call Progress Tone Generator Package
 - Special Information Tone
 - Ringing Tone
- H.248.27 Conferencing Tones Generation Package
 - Conference Enter Tone
 - Conference Exit Tone

The played tone always temporarily interrupts the speech.



The Tone Sender service can be configured through the TsTone MO. Configurable tone parameters include tone type, tone duration, frequencies, levels, play and pause times. For more information, see [vMRF Configuration Management](#).

13.14 User Plane Frame Handler Service (UP FH)

The User Plane Frame Handler (UP FH) service provides RTP payload framing according to the corresponding RTP profile.

14 Support for Watchdog

The watchdog feature, when configured, supervises the status of each VM instance and restarts the instance if the watchdog does not get signal from the VM for one minute. This time period is a hard-coded value in the driver of the watchdog device, so it cannot be configured.

The watchdog device consists of an emulated Intel 6300 ESB device running within the hypervisor and a watchdog daemon running on each VM instance.

To use the watchdog feature, the following conditions must be met:

- The appropriate software delivery package that includes watchdog device drivers and daemons is used.
- The VNF flavor stack used contains the watchdog action.
- Watchdog deployment parameters are configured in the `example_environment.yaml` file during deployment time.

Note: The watchdog device is pinged only when the flavor is configured with the watchdog action as well, otherwise changing watchdog parameter values has no result.

Besides periodically signaling the watchdog device in the hypervisor, the watchdog daemon also supervises certain system resources in the VM, and can forcefully restart the VM instance by itself in the following cases:

- The watchdog daemon cannot open the `/proc/uptime` file.
- The kernel process table is full (that is, the watchdog cannot create new processes).

The watchdog feature has the following configurable parameters:

- Ping interval
- Parameter for enabling or disabling watchdog

For default values, see the relevant [deployment documentation](#).

For changing to non-default parameter values (including enabling or disabling the feature), the `example_environment.yaml` HOT file must be updated during deployment time.

While a monitored system parameter indicates fatal state, no heartbeats are sent to the watchdog device, so it performs the configured watchdog action (as set during flavor creation).



15 VMware Tools

vMRF includes Open VM Tools (OVT) in the guest OS. OVT is an open source implementation of VMware Tools. It consists of virtualization utilities that improve the functionality and management of VMs within a VMware environment.

[Table 1](#) describes how supports OVT.

Table 1 OVT Support in vMRF

OVT Feature	Support in vMRF
Synchronize the guest OS clock with the virtual infrastructure.	Not to be used. vMRF synchronizes the clock with an external NTP server.
Enable the virtual infrastructure to perform graceful shutdown of the VM.	Supported. The <code>administrativeState</code> attribute of the <code>MrfInstance</code> MO representing the VM remains UNLOCKED while the <code>operationalState</code> attribute is set to disabled and the <code>availabilityStatus</code> attribute is set to POWER_OFF.
Enable the virtual infrastructure to quiesce the file system of the VM before taking a VM snapshot.	Quiescing the file system of the VM before taking a VM snapshot is not supported.
Provide a heartbeat from the guest to the virtual infrastructure to support vSphere High Availability (HA).	VM monitoring supported.
Publish information about the guest OS to the virtual infrastructure, including resource utilization and networking information.	IP addressing information supported.
Provide a secure and authenticated mechanism to perform various operations within the guest OS from the virtual infrastructure.	Not supported. The standard vMRF O&M interfaces provide accountability logging and enforcement of role-based access rights on VNF level and must therefore be used.
Customize the guest OS when deploying a new VM.	Not supported. For more information about configuration injection in vMRF, see to the relevant deployment instructions .
Improve the interaction with the guest desktop environment.	Not supported. vMRF does not provide a desktop environment.



16 Limitations and Differences

The following MRF functionality is currently not supported in vMRF:

- MRFC functionality, that is, Mr interface (NetAnn, MSCML, and MSML)
- Video streams (video codecs H.264 and VP8)
- Audio codecs EFR, EVRC, EVRC-B, EVRC-NW, Opus, and G.719
- External announcements stored to a web server (fetched using HTTP)

Normally vMRF ignores and discards, or rejects H.248 packages and properties related to the unsupported features. To avoid confusion and unnecessary troubleshooting, all H.248-controlled features that are not supported by vMRF must be turned off in the MTAS configuration.

The following functionality works differently in MRF and vMRF:

- The supported encoding of audio announcements has been changed from raw format 8 bit A-law or μ -law PCM in MRF to WAV format 16 bit 16 kHz linear PCM in vMRF.
- The handling of variable announcements has been changed. The algorithm and configuration files for variable announcements in MRF have been replaced by lua scripts in vMRF.

16.1 vMRF VNF Size Limitations

The minimum and maximum number of VMs in a vMRF cluster and the minimum and maximum number of vCPUs per VM are shown in the table below.

Deployment Scenario	Number of VMs per VNF	Number of vCPUs per VM
vMRF	2–20	4–52

The maximum number of vCPUs per VM is also limited by the number of physical CPU cores in the underlying hardware architecture. This means that, depending on the infrastructure, the maximum number of VMs and vCPUs can be lower than the software limitations described in the table.



Reference List

7 kHz audio-coding within 64 kbit/s, [ITU-T G.722](#)

G.711: Pulse code modulation (PCM) of voice frequencies, [ITU-T G.711](#)

G.729: Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear prediction (CS-ACELP), [ITU-T G.729](#)

IPv6 Stateless Address Autoconfiguration, [IETF RFC 4862](#)

RTP: A Transport Protocol for Real-Time Applications, [IETF RFC 3550](#)

RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals, [IETF RFC 4733](#)