



Music on Hold with IP Connectivity

And an introduction to Network Regions

Abstract

Music On Hold has long been a standard feature in Avaya communications systems. Traditional TDM based systems take a music source and interconnect it as needed throughout the system on time slots. This is done in a very efficient manner that allows hundreds of endpoints spread across many port networks to be connected to the music source without using up multiple instances of critical system resources. The IP enabled systems have some new concepts to be aware of when implementing Music On Hold.

This paper provides a description of how Music On Hold operates and provides instructions on how it can be implemented using the Network Regions feature of Avaya Communication Manager

1.	Introduction.....	3
2.	Basic Circuit Switched Music On Hold.....	3
3.	Adding IP - Music on Hold to a Media Gateway	4
4.	Adding IP stations to the mix.....	4
5.	Adequate bandwidth - Codecs and Music	5
5.1.	Remote Facilities – WAN connectivity	6
6.	Problem – Different configurations have different needs.....	7
7.	Network Regions	7
8.	Codec Sets.....	9
9.	How do endpoints get associated with network regions?	10
10.	Multiple resources in a PN – which one to use?	13
10.1.	What if a resource isn't available?	14
11.	What about call centers?	14
12.	COR Override – remote IP stations	16
13.	Multiple Music Sources – Tenant Partitioning	16
14.	Related Features.....	16
15.	Alternative.....	16
16.	Conclusion	16
17.	Additional References.....	16

1. Introduction

Music On Hold (MOH) has long been a standard feature in Avaya systems. In the traditional Time Domain Multiplexer (TDM) based systems a music source is interconnected as needed throughout the system on time slots. This is done in a very efficient manner that allows hundreds of endpoints spread across many port networks to be simultaneously connected to the music source without using up multiple instances of critical resources. The IP enabled systems are not quite as efficient and have some new concepts to be aware of when implementing music on hold. This paper provides a description of how MOH operates and provides instructions on how it can be implemented using the Network Regions feature of Avaya Communication Manager.

Note: This document contains a tutorial to explain how network regions and codec sets relate to the music-on-hold feature. This is only one aspect of the subject of network regions. Additional important information about network regions is available in references one and two.

2. Basic Circuit Switched Music On Hold

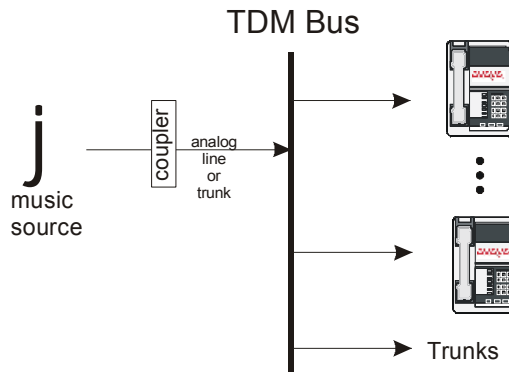


Figure 1: Basic Circuit Switched Music On Hold

Figure 1 shows the most fundamental case. The music source is connected through an audio coupler to an analog line. It gets automatically connected to a timeslot when music is needed anywhere in the system. This timeslot can be listened to by as many endpoints, stations or trunks, as desired. The timeslot is released when there is no need for MOH anywhere in the system.

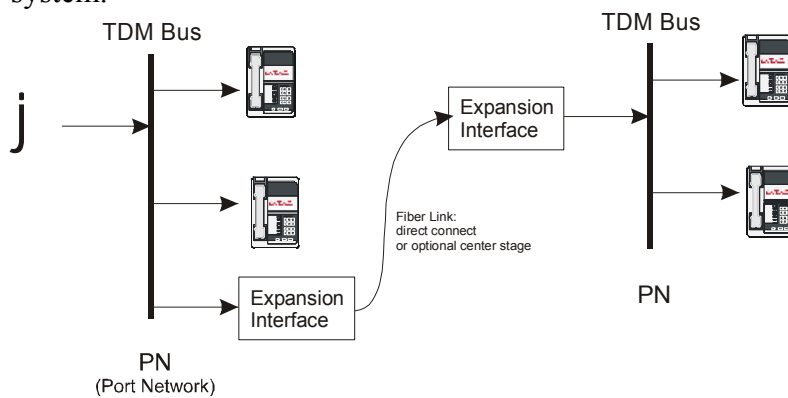


Figure 2: Basic Circuit Switched Music On Hold with multiple port networks

Figure 2 shows how this feature extends to multiple port networks. Note that only a single timeslot is used in each port network and on the fiber link serving any endpoint that needs MOH. This feature scales up very well.

3. Adding IP - Music on Hold to a Media Gateway

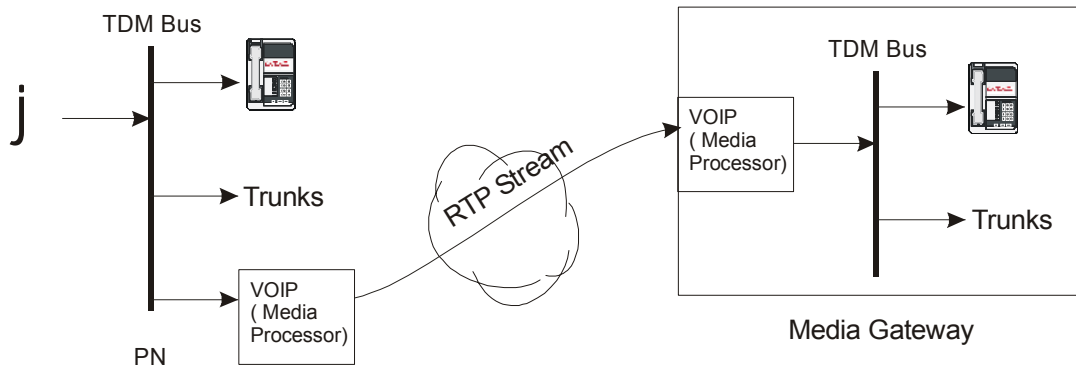


Figure 3: Basic Packet Switched Music On Hold with multiple port networks

When a Media Gateway is added to the configuration the same strategy is followed. In this case, instead of a time slot over the fiber, a VOIP Media Processor packages the Pulse Code Modulation (PCM) samples into IP packets which are transported as a Real Time Protocol (RTP) stream through the network to the gateway. Another Media Processor in the gateway converts incoming packets back to PCM and delivers the PCM to the TDM bus. Like before, many endpoints on each gateway, be they stations or trunks, can listen to the same stream at the same time.

It is important to note that the music source can be in either a Port Network (PN)¹ or a Media Gateway (MG) and the transmission between functional units works equally well in either direction. Also, this works the same way and scales well with port network groups (port networks controlled by a single IPSI and interconnected by TN570x Expansion Interfaces), even if one PN has no Media Processors. The MOH transmission path would be the media processor, the local TDM bus, the EI fiber link and the TDM bus on the other PN.

4. Adding IP stations to the mix

The above systems only consider endpoints served directly by the gateway TDM busses or PNs. The situation gets more complicated as IP stations and remote sites are included. As shown below, each remote IP endpoint goes directly back to the music source for MOH.

¹ Although both are true gateways, for this document a GW with TN coded circuit packs is referred to as a port network (PN).

This is fine if all the endpoints are on a LAN. If there is a WAN link involved, it's important to provide adequate bandwidth for all IP endpoints that will be listening to MOH and are remote over the WAN connection.

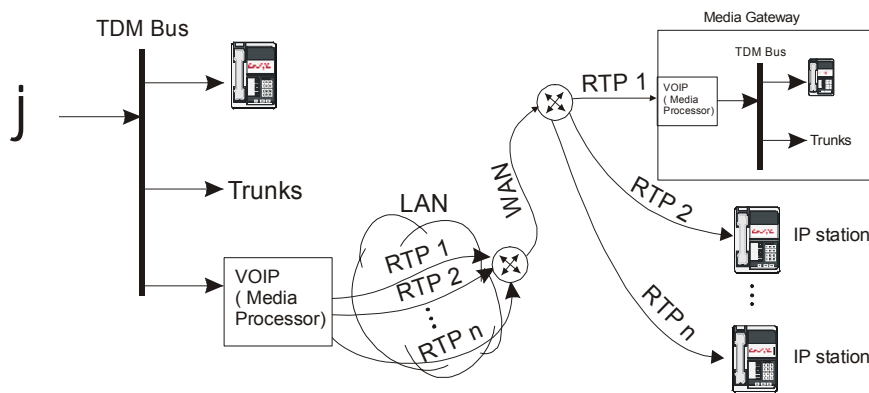


Figure 4: Each IP endpoint gets its own RTP stream

5. Adequate bandwidth - Codecs and Music

The media processors shown in the previous figures use a function called a coder-decoder (codec) to take standard Pulse Code Modulation (PCM) streams, packetize and transmit them using the Real Time Protocol (RTP) stream across the network. The same codec function on the receiving end of the stream disassembles the packets and re-creates the PCM stream. There are different kinds of codecs that vary from fairly straightforward raw delivery of PCM streams (G.711) to compression schemes used to conserve bandwidth. Common compression schemes used today can compress the normal 64kbps PCM streams to around 8kbps or less (ignoring packet overhead). They are highly optimized towards human voice, not music. G.711 typically delivers the best quality music-on-hold whereas the compressed codecs (G.729, G.729a, G.729b, G.729AB, G.723, etc) don't sound quite as good and many times are not acceptable for MOH. For the best results, when possible use G.711 as the codec of choice for delivering MOH. Although G.711 uses more bandwidth on the network than the compressed codecs, it is simple and straightforward and common for endpoints on a LAN where bandwidth is plentiful and resources are relatively inexpensive.

The following figure shows a high capacity LAN spanning a campus installation. There is plenty of capacity between sites and there is no reason to use compressing codecs to save on bandwidth.

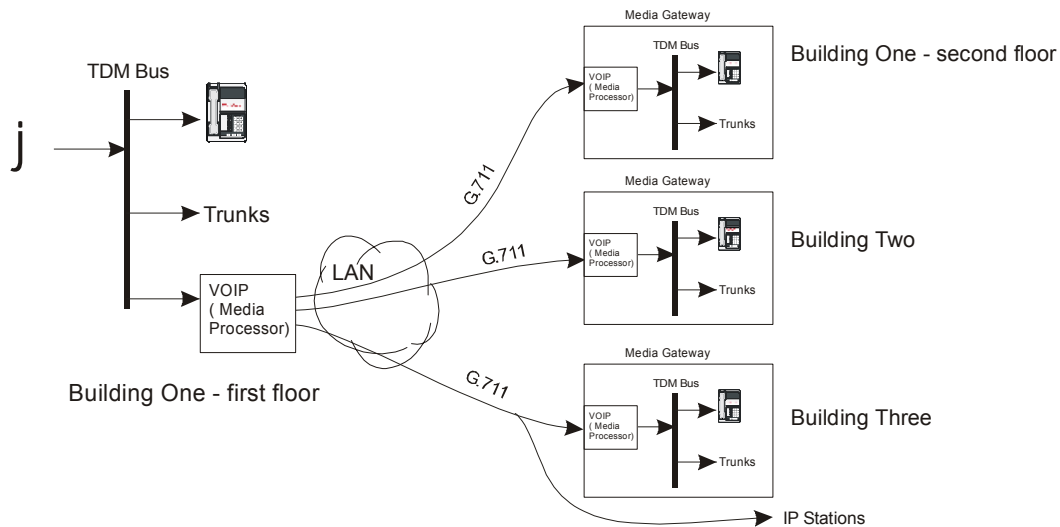


Figure 5: A campus environment

5.1. Remote Facilities – WAN connectivity

Connections to remote facilities require bandwidth which is typically more expensive. Many are willing to accept some degradation in voice quality to and from those specific facilities in exchange for the bandwidth savings. Although fine for voice, many feel the music quality is unacceptable using the compressed codecs. For example, Figure 6 shows the same configuration as above but buildings 2 and 3 are now remote from building one and therefore use more expensive bandwidth. Avaya Communication Manager tries to be efficient and, no matter how many endpoints on a GW need MOH, it still only creates one stream of music across the WAN to the GW. This single stream services all endpoints on that GW and only exists when one or more need to hear music, otherwise the RTP stream is discontinued and the timeslot is released. Using G.711 would be acceptable to get good quality MOH but could become expensive to carry all of the voice traffic to and from the remote sites.

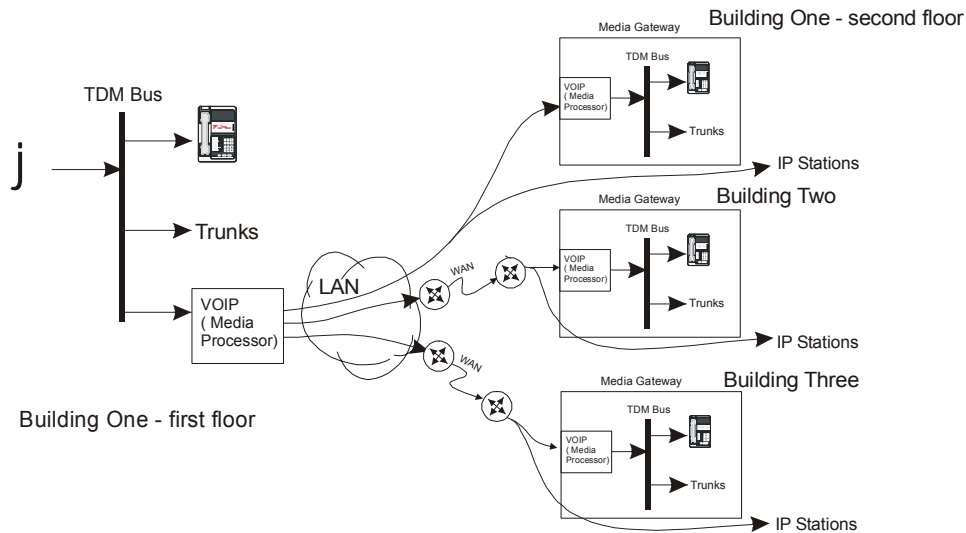


Figure 6. Music across the WAN

6. Problem – Different configurations have different needs

Systems have different transmission needs depending on if they are transporting voice or music, what endpoints are connected together, and what network resources are available between those endpoints. There is no way to effectively distinguish music from voice in a connection. A mechanism is needed to choose different codecs for different cases.

Avaya Communication Manager uses network regions to help specify which codec to use for IP connections between two endpoints. Think of a network region as a geographic area of high bandwidth availability. These IP network regions are generally interconnected to other regions by lower bandwidth WAN facilities. Avaya Communication Manager allows specification of codec-sets to optimize for voice quality for some connections and select low bandwidth for others. As described above, higher voice quality (music) can be specified, but it comes at a price of higher bandwidth usage.

7. Network Regions

Avaya Communication Manager provides the ability to identify and assign endpoints, resources, etc to different groups called Network Regions. It allows specification of parameters such as which codec to use for transmission within each region, but more importantly it allows specification of which codecs and parameters to use when creating connections between regions. Figure 7 shows a system with the different locations divided into network regions. Each region was chosen to be a group of resources and endpoints that share a common set of characteristics. The important resource in region one is the music source which all other regions need access to. Region two has a geographical characteristic which shares a high speed LAN with region one. Regions 3 and 4 are remoted over WAN connections where it is desirable to minimize bandwidth usage.

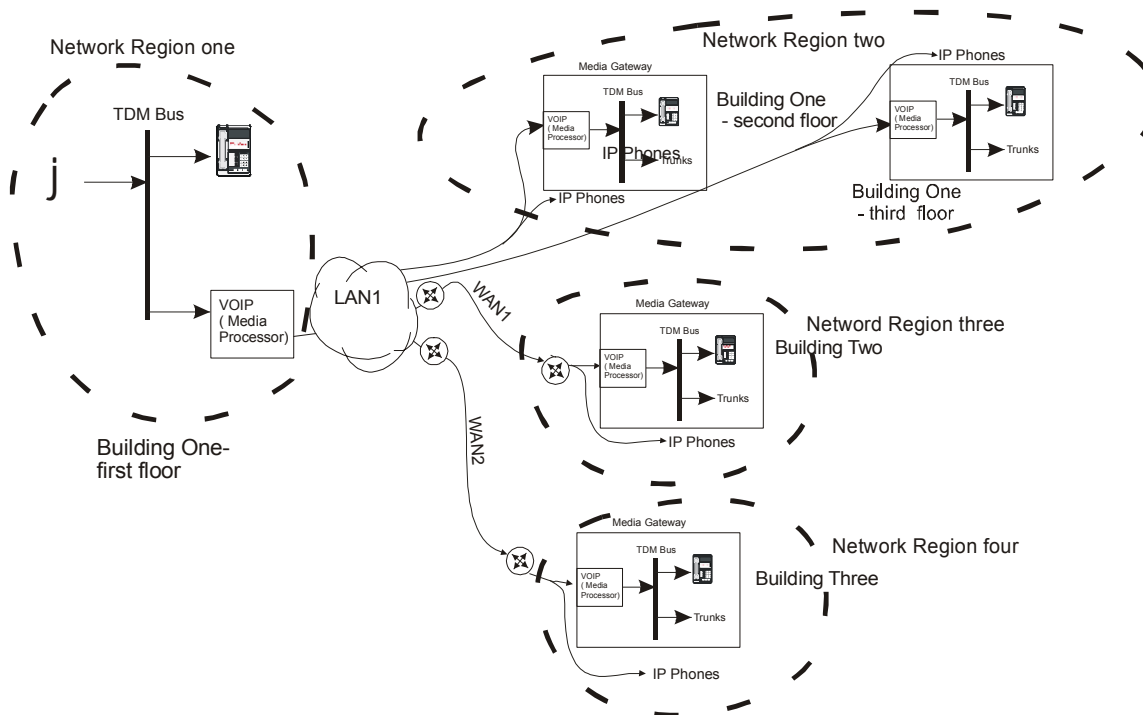


Figure 7. Dividing a system into Network Regions

All Network Regions, for this example, have a local LAN and would be set up to use G.711 within the region. This is controlled by specifying what is called a codec-set on the first page of the ip-network-region form. In this case codec-set one.

```

change ip-network-region 1                                     Page 1 of 3
IP NETWORK REGION

Region: 1
Location:
AUDIO PARAMETERS
Codec Set: 1
UDP Port Range
Min: 2048
Max: 3028

DIFFSERV/TOS PARAMETERS
Call Control PHB Value: 34
Audio PHB Value: 46

802.1P/Q PARAMETERS
Call Control 802.1p Priority: 7
Audio 802.1p Priority: 6

Name:
Intra-region IP-IP Direct Audio: yes
Inter-region IP-IP Direct Audio: yes
IP Audio Hairpinning? y

RTCP Reporting Enabled? y
RTCP MONITOR SERVER PARAMETERS
Use Default Server Parameters? y

AUDIO RESOURCE RESERVATION PARAMETERS
RSUP Enabled? n

ESC-x=Cancel Esc-e=Submit Esc-p=Prev Pg Esc-n=Next Pg Esc-h=Help Esc-r=Refresh

```

The IP Network Region form specifies more than just what codec-set to use, it specifies many of the characteristics of the transmission including UDP port range, QOS parameters, IP-IP Direct capabilities, etc.

8. Codec Sets

The codec-sets define Codecs in order of preference and packet sizes. Seven codec-sets may be defined, more than enough to define the different needs within a specific system. There is a preference list of codecs in a codec-set to allow endpoints with different capabilities to be served efficiently. Simply speaking, when setting up a connection the two endpoints start at the first preference and work their way down the list until they find a codec setting they both can do. In the following example, codec-set one has the first (and only) choice of 20msec packets of G.711. No other choices are needed because it is known that all endpoints in NR (network region) one can support this choice.

```
display ip-codec-set 1
```

```

IP Codec Set

Codec Set: 1

Audio      Silence      Frames      Packet
Codec      Suppression  Per Pkt     Size(ms)
1: G.711MU          n           2           20
2:
3:
4:
5:
6:
7:

```

```
Command:
```

```
ESC-x=Cancel Esc-e=Submit Esc-p=Prev Pg Esc-n=Next Pg Esc-h=Help Esc-r=Refresh
```

A very important part of network regions is allowing connections between regions independently of the intra-region setting. These connections would be controlled by the network region connectivity matrix (page 2 of the NR forms) shown below.²

```
change ip-network-region 1
```

```
Page 2 of 3
```

```

Inter Network Region Connection Management

Region
          (Group Of 32)
1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2
001-032 1 1 1 1
033-064
065-096
097-128
129-160
161-192
193-224
225-250

```

² This form will change starting in Avaya Communication Manger release 2.0 to introduce additional capabilities.

The left most column defines a range of regions that can be interconnected to this region. Each position to the right is for a specific NR in that range. Each of the 250 possible destination regions has a position in this table. For each intersection, enter the number of the codec-set to use when making a connection between an IP interface in this region and one in the destination region. Referring back to figure 7, to get all regions to use G.711 when connecting to the music source, specify codec-set one between NR1 and all other NRs.

Important: Specify G.711 between NR 1 and all other NRs to get good music quality to anywhere in the system. This also chooses G.711, with its higher bandwidth requirements for all other connections to NR 1. Make sure the required bandwidth is provided for these connections.

The following screen for NR3 shows that the inter-network region connection capability can specify building 2 (region 3) use G.711, but from region 3 to the other buildings use G.729 (codec set 2). It works the same for building 3 (region 4).

```
change ip-network-region 3 Page 2 of 3
Inter Network Region Connection Management
Region <Group Of 32>
001-032 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2
033-064 1 2 1 2
065-096
097-128
129-160
161-192
193-224
225-250
```

Building one, second and third floors are in the same region and use G.711 because they share a LAN and don't need to be as concerned about bandwidth. This example would define codec set two as G.729 to be used for inter-region connections.

```
change ip-codec-set 2 Page 1 of 1 SPE A
IP Codec Set
Codec Set: 2
Audio Silence Frames Packet
Codec Suppression Per Pkt Size(ms)
1: G.729A n 2 20
2: _____ - -
3: _____ - -
4: _____ - -
5: _____ - -
6: _____ - -
7:  _____ - -
```

9. How do endpoints get associated with network regions?

As described above, network regions are used to select a codec set to be used for the connection between any two IP interfaces on a call-by-call basis. Therefore, all IP interfaces must be assigned a NR.

IP Stations – IP stations can be assigned to a network region in two ways.

IP INTERFACES

ON	Type	Slot	Code	Sfx	Node Name	Subnet Mask	Gateway	Address	Net Rgn	ULAN
y	C-LAN	01A04	TN799	D	Clan1	255.255.255.0	172.19	.27 .220	3	n
y	MEDPRO	01A10	TN2302		Prowler1	255.255.255.0	172.19	.27 .220	3	n
y	C-LAN	02A04	TN799	D	Clan2	255.255.255.0	172.19	.18 .254	1	0
y	MEDPRO	02A09	TN2302		Prowler2	255.255.255.0	172.19	.18 .254	1	0
y	MEDPRO	02A10	TN2302		Prowler3	255.255.255.0	172.19	.18 .254	1	0
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n
n						255.255.255.0	.	.	.	n

ESC-x=Cancel Esc-e=Submit Esc-p=Prev Pg Esc-n=Next Pg Esc-h=Help Esc-r=Refresh

Media Gateways -The NR is specified on screen one of the Media Gateway administration form.

IP Trunks - IP trunks, like all IP connections, have a near-end region and a far-end region. The network region for the near end is the same as the CLAN specified on the signaling-group form by the "Near-end Node Name". The far end region is assumed to be the same as the near end unless specified to be different by the "Far-end Network Region" field. Like any other IP connection, a connection across an IP trunk chooses a codec-set specified by the intersection of these network regions.

SAT satur

display signaling-group 7 Page 1 of 5

SIGNALING GROUP

Group Number: 7 Group Type: h.323
 Remote Office? n Max number of NCA TSC: 0
 SBS? n Max number of CA TSC: 0
 Trunk Group for NCA TSC: 7

Trunk Group for Channel Selection: 7
 Supplementary Service Protocol: a

Near-end Node Name: satclan2 Far-end Node Name: AVAYA-CLAN
 Near-end Listen Port: 1720 Far-end Listen Port: 1720
 Far-end Network Region:

LRQ Required? n Calls Share IP Signaling Connection? y
 RRQ Required? n Bypass IF IP Threshold Exceeded? y
 Media Encryption? n

DTMF over IP: out-of-band Direct IP-IP Audio Connections? n
 IP Audio Hairpinning? n
 Interworking Message: PROGRESS

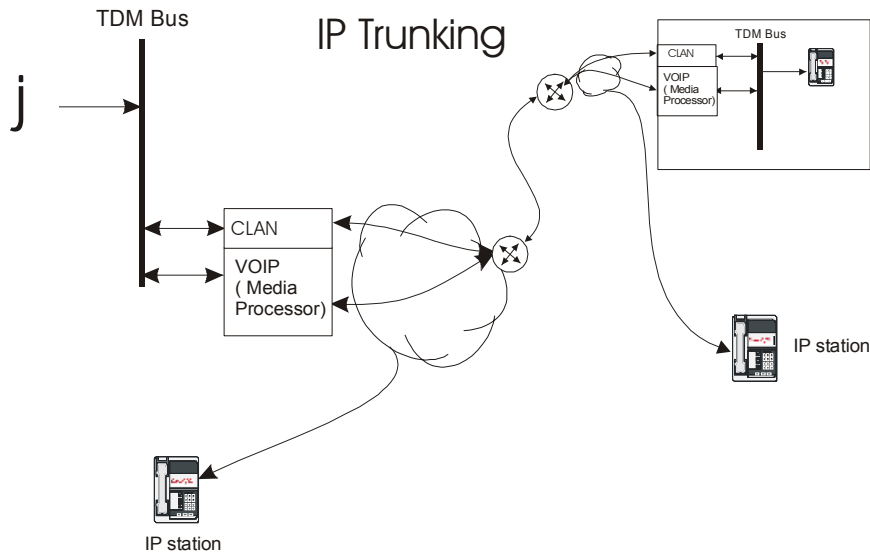


Figure 8: IP Trunks

A connection between the MOH source and an IP trunk is treated as a connection between the Media Processor serving the MOH source and the far-end network region.

10. Multiple resources in a PN – which one to use?

An additional way codecs are chosen is by preferring same-region connectivity. For example, Fig 9 shows a PN with media processors and CLANs in NRs 3, 4, and 6. There are also GWs and IP stations in NRs 3, 4, and 6.

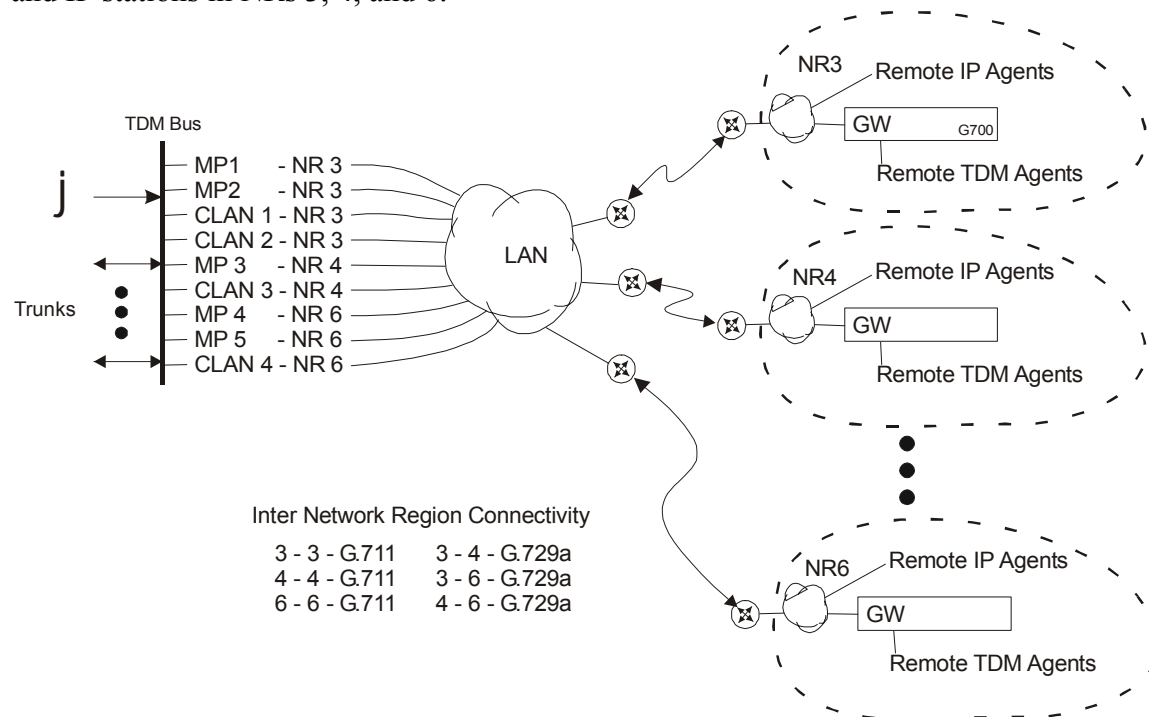


Figure 9: Multiple Resources

When making connections, Avaya Communication Manager will always choose devices in the same network region first. So a connection from an IP phone in NR3 to a trunk or the music source will be steered to a media processor in NR3 (MP1 or MP2) and use G.711 (3-3). Likewise, a connection to a trunk or the music source from NR4 would use MP3 and G.711 if available. A connection to a trunk from an IP station in NR4 will prefer a NR4 media processor and use G.711 because it's a 4-4 connection.

10.1. What if a resource isn't available?

If there are no Medpro resources in the same NR available, rather than reject the connection it will choose equivalent devices not in the same region, but only if they are connected via the inter-network-region connection form.

For example, for a trunk connection to NR3 and all resources in MP1 and MP2 are in use AND network region 3 is connected to network region 4 on the inter-network-region form, it will use a MP in region 4 and it will use the codec set for an inter network region connection between NR3 and NR4 (G.729a).

11. What about call centers?

What about call centers where there are large numbers of incoming (or outgoing) trunk calls that can be distributed to remote agents? The agents don't need to listen to MOH, the trunks do. So the goal is to keep the MOH source as near to the trunks as possible and provide adequate bandwidth. Many call centers aggregate their trunk facilities in one or two locations and have agents both local to the trunk facilities and remote. A common scheme is shown in figure 10.

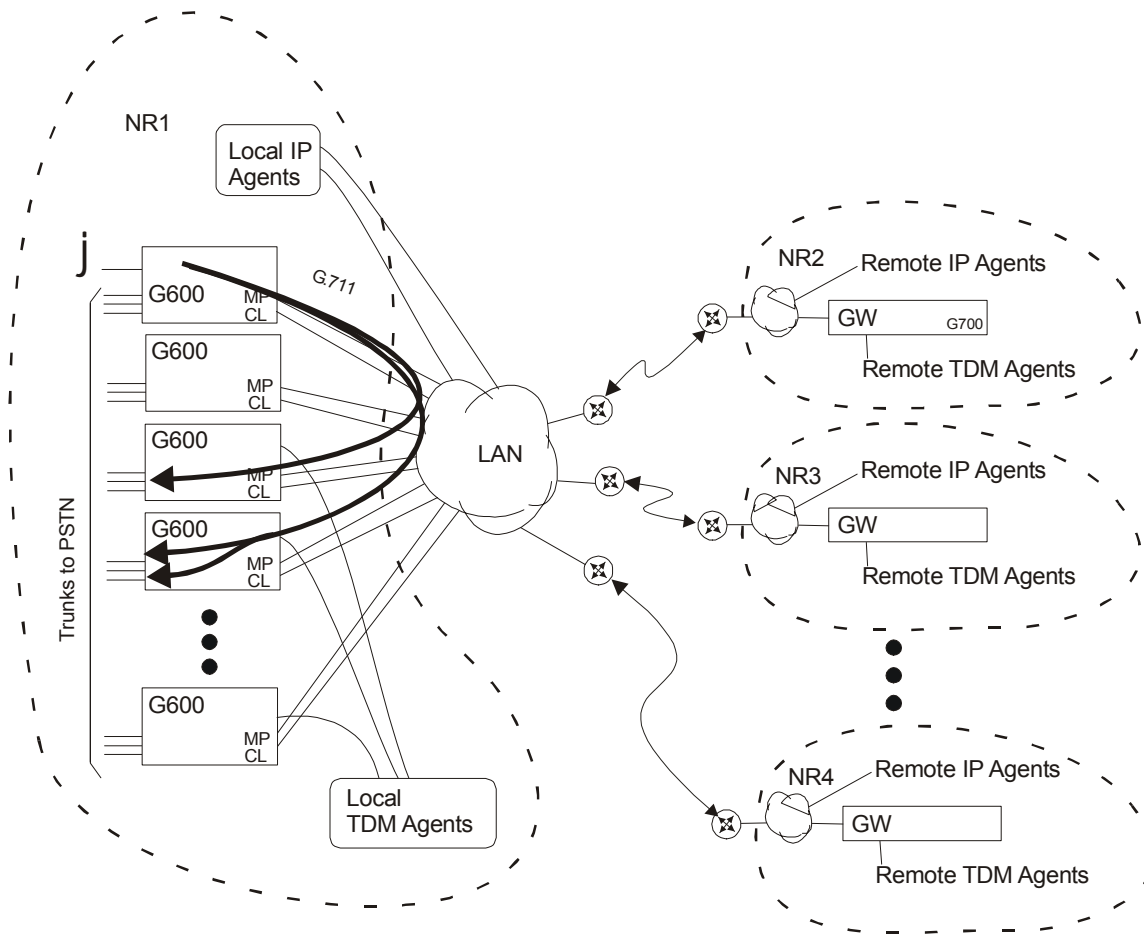


Figure 10: Typical call center – multiple trunks receiving MOH

Heavy call center traffic on the trunks is distributed as desired between agents local and remote. The local trunks, some agents, and MOH source are in NR one. Each remote location is in its own NR so it can be independently controlled. All NRs are set to use G.711 internally and use G.729a or b to other NRs. Because all the trunks and the music source are in NR 1, when a customer on a trunk port is put on hold, the trunk (NR1) is connected to the music source (NR1) and G.711 is used. Trunk connections (NR1) to remote agents (NR2, 3,4) use G.729x to conserve bandwidth. In this case, for connections between agents and the trunks, the agents must have their COR set to "hear music on hold=y" otherwise when they put a call on hold the trunk caller would not hear MOH.

Agents rarely get put on hold, so they don't clog the WAN connections with G.711 connections to the NR1 music source.

12. COR Override – remote IP stations

The IP stations in remote sites could use more bandwidth over the WAN than initially realized since each would get its own separate connection to the MOH source. It's not a major issue because the agents rarely get put on hold, they normally put trunk calls on hold. The COR setting of the station where the HOLD button is pushed is used to decide if the other party hears MOH or not. So, as noted above, if an agent wants to press the HOLD button and have the other party hear MOH, then that agent's COR must have "Hear Music on Hold" set to yes. There may be some instances and configurations where it would be desirable to save bandwidth by not using MOH on certain stations.

This behavior is unchanged from previous versions of the Avaya products and independent of any IP functionality.

13. Multiple Music Sources – Tenant Partitioning

These examples have only utilized the single System Music Source. The Tenant Partitioning feature (system-parameters features screen) allows one music source per partition. Network regions work the same way within Tenant Partitioning.

14. Related Features

Call Parking, First Announcement Delay, Hospitality/auto-wakeup

15. Alternative

If, for some reason, a system can't be configured to meet the design goals, consider a 'voice-on-hold' source, like a news channel or a repeating voice announcement. These are non-music sources that work well with the lower bandwidth codecs.

16. Conclusion

This paper describes how Music On Hold is implemented in Avaya Communication Manager systems. It points out that common IP codec compression schemes are optimized to human voice, not music, and therefore that Music On Hold works best with G.711 uncompressed codec. It shows how codecs are selected using the Network Regions feature and provides an introduction to the use of the network regions feature. It describes issues involved to design your system to best provide Music On Hold for both Media Gateways and IP stations.

17. Additional References

- 1) Network Regions For Avaya MultiVantage™ Solutions – A Tutorial
- 2) Avaya IP Telephony Implementation Guide – May 2003
- 3) Administration Guide for Avaya™ Communication Manager – 555-233-506
- 4) 4600 Series IP Telephony LAN Administrators Guide

Notice: While reasonable efforts were made to insure that the information in this document was complete and accurate at the time of printing, Avaya Inc. can assume no liability for any errors. Changes and corrections to the information in this document may be incorporated in future releases.

©2003 Avaya Inc. All Rights Reserved.

Avaya and the Avaya Logo are trademarks of Avaya Inc. All trademarks identified by ® and ™ are registered trademarks or trademarks, respectively, of Avaya Inc. All other trademarks are the property of their respective owners. The information provided in these Application Notes is subject to change without notice. The configurations, technical data, and recommendations provided in these Application Notes are believed to be accurate and dependable, but are presented without express or implied warranty. Users are responsible for their application of any products specified in these Application Notes.

Please e-mail any questions or comments pertaining to these Application Notes along with the full title name and filename, located in the lower right corner, directly to the Avaya Solution & Interoperability Test Lab at interoplabnotes@list.avaya.com