Video Services

In This Section

This section describes how to configure the hardware for video services and some basic video services configuration concepts in support of the IPTV video applications.

Topics include:

Video Services

Video Groups

When configured in the router, ISA-MS are logically grouped into video groups for video services. A video group allows more than one video ISA to be treated as a single logical entity for a given application where the system performs a load balancing function when assigning tasks to a member of the group. All video group members are “active” members, so there is no concept of a “standby” ISA as in other ISA groups in the 7750 SR and 7450 ESS.

Video groups provide a redundancy mechanism to guard against hardware failure within a group where the system will automatically rebalance tasks to the group excluding the failed ISA. Video groups also pool the processing capacity of all the group members and will increase the application throughput because of the increased packet processing capability of the group. The buffer usage is typically identical for all members of the video group, so increasing the number of members in a group will not increase the scaling numbers for parameters bounded by available buffering, but there will still be the increase in performance gained from the pooled packet processor capacity. A video service must be enabled at the video group level before that service can be used.

A maximum of four ISA-MSs can be supported in a single video group. Note that a given video application may restrict the number of members supported in a video group to a smaller number. Refer to specific sections in this guide for video application additional information.

A maximum of four video groups are supported in a router. There is a chassis limit of eight ISA-MSs per router which constrains the number and members of video groups.

Note: ISA-MS in a single video group cannot be on the same IOM. An IOM can accomodate two ISA-MS modules provided that the ISA-MS are members of different video groups.

Video SAP

The video group logically interfaces to a service instance with a video Service Access Point (SAP). Like a SAP for connectivity services, the video SAP allows the assignment of an ingress and egress filter policy and QoS policy.

Note: Ingress and egress directions for the filter and QoS policy are named based on the perspective of the router which is the opposite perspective of the ISA. An “egress” policy is one that applies to traffic egressing the router and ingressing the ISA. An “ingress” policy is one that applies to traffic ingressing the router and egressing the video. Although potentially confusing, the labeling of ingress and egress for the ISA policies was chosen so that existing policies for connectivity services can be reused on the ISA unchanged.

If no filter or QoS policy is configured, the default policies are used.

One of the key attributes of a video SAP is a video group association. The video SAP’s video group assignment is what determines which video group will service on that video SAP. The video groups configuration determines what video services are available.

Video Interface

A video interface is a logical IP interface associated with a video SAP and provide the IP addressing for a video SAP.

A video interface can have up to 16 IP addresses assigned in a Layer 3 service instance. A video interface can have only one IP address assigned in a Layer 2 service instance.

Multicast Information Policies

Multicast information policies on the 7750 SR and 7450 ESS serve multiple purposes. In the context of a service with video services, the multicast information policy assigned to the service provides configuration information for the multicast channels and defines video policy elements for a video interface.

Note: This section describes the base elements of a multicast information policy in support of a video service. Specific video service features will require additional configuration in the multicast information policy which are described in the sections dedicated to the video feature.

Multicast information policies are named hierarchically structured policies composed of channel bundles which contain channels which contain source-overrides.

- Bundles are assigned a name and contain a collection of channels. Attributes not defined for a named bundle are inherited from the special default bundle named “default”.

```
*A:ALA-48configmcast-mgmtmcast-info-plcy# info
----------------------------------------------
            bundle "default" create
            exit
----------------------------------------------
*A:ALA-48configmcast-mgmtmcast-info-plcy#
```

- Channels are ranges of IP multicast address identified by a start IP multicast address (GG_{start}) and optional end IP multicast address (G_{end}), so the channels encompasses (*,G_{start}) through (*,G_{end}). A channel attribute is inherited from its bundle unless the attribute is explicitly assigned in which case the channel attribute takes precedence.
- A source-override within a channel are IP multicast addresses within the channel with a specific source IP address ($S_{override}$), so the source-override encompass ($S_{override}$,G_{start}) through ($S_{override}$,G_{end}). A source-override attribute is inherited from its channel unless the attribute is explicitly assigned in the source-override channel in which case the source-override channel attribute takes precedence.

For a given IP multicast channel (*,G) or (S,G), the most specific policy element in the hierarchy that matches applies to that channel.

A multicast information policy is assigned to a service instance. For video services, the multicast information policy assigned to the service determines the video group for a given IP multicast channel. When a channel is assigned to a video group, the channel is sent to the video group for buffering and/or processing as appropriate depending on the video services enabled on the video group. If no video group is assigned to a given channel, the channel will still be distributed within the service instance, but no video services will be available for that channel.

In addition to bundles, channels and source-overrides, multicast information policies also include video policies. Video policies define attributes for the video interfaces within the service instance.

Note: Video policy attributes are specific to the video feature and will be covered in detail in the applicable video feature section. Video policies are mentioned here because they are an element of the multicast information policy and provide the link to configuration for a video interface.

Duplicate Stream Protection

While H-RET can protect against minor amounts of packet loss, it is limited in the number of packets that can be recovered (currently 32). This can be from approximately 125ms of a 3Mbps stream to only 18ms for a 20Mbps stream. These times are short for a network reconvergence event which will typically be in the order of 300-1200ms. Further, retransmission will cause incremental bandwidth spikes in the network as the lost packets are sent to the client as quickly as possible.

Rather than invoke a retransmission event to protect against network interruption or reconvergence, it is often more efficient to protect the stream via an alternate transmission path. This can be a separate physical interface, transmission link, system or even technology.

Duplicate-stream protection allows an operator to split a single multicast stream (single S,G and common SSRC) into two different transmission paths that may have different transmission characteristics (latency/jitter). Rather than select one stream for retransmission to the client the Duplicate Stream protection feature evaluates each stream packet-by-packet, selecting the packet that first arrives (and is valid) for retransmission.

A circular buffer is used for duplicate-stream protection which incorporates both packet-by-packet selection (based on RTP sequence number/timestamp and SSRC) and a re-ordering function whereby any out-of-sequence packets will be placed into the buffer in order, thus creating a corrected, in-order stream.

Similar to the H-RET re-sequencing feature playout rate is a function in ingest rate, however because the two streams may be delayed between one-another a few assumptions are made:

- The first arriving packet is always put into the buffer, allowing for the backup medium to wander in terms of latency and jitter.
- Because the source is the same, the rate at which a packet is put into the buffer (from either stream) can be assumed to be the normal bitrate.

The output RTP stream is always maintained in-sequence and the playout speed is user-controlled. Either with constant-delay (i.e., packet ingress time + 500ms = packet egress time) or can be a moving window average to smooth jitter that may occur between packets or the two contributing streams. The operator can specify the size of this window where zero (0) is a constant-delay.

The buffer size is similarly configurable and is the higher of the inter-stream phase (i.e., one stream ahead of another) or the expected jitter.

Duplicate Stream Selection

Stream Identification

Stream selection is a simple selection algorithm that is applicable to any number of input streams. It is a prerequisite for stream selection that RTPv2 encapsulation be used in UDP.

Each service is identified by multicast source, group/destination address and current synchronization source (SSRC). Once this has been identified, the ISA monitors its ingress for:

- Traffic with a DA of the multicast group, or;
- Traffic with a DA of the ISA (unicast)

Traffic is further checked as having RTP-in-UDP payload, RTP version 2.

The SSRC of each incoming RTP packet is learned as unique sources. Only one SSRC is supported for each stream, however as SSRC may change during abnormal situations (such as encoder failover), it can be updated.

A SSRC can only be updated when a Loss of Transport (LoT) occurs, as other duplicate streams (with the original SSRC) may still be operational. When an LoT occurs the SSRC is deleted, the buffers are purged and the RTP sequence counters are reset. The SSRC will be extracted from the next valid RTP packet and the sequence will start over.

Note that individual streams are not tracked by the ISA. There may be one, two, or ten duplicate streams, the number is of no consequence to the selection algorithm (however bandwidth and/or video quality monitoring (VQM) may be impacted). Irrespective of the number of duplicate streams, one RTP packet is selected for insertion into the video ISA buffer. Once a packet is selected the RTP sequence counter is incremented and any further RTP packets received by the ISA with the previous sequence number are discarded.

In summary, duplicate stream selection is a FIFO algorithm for RTP packet selection, this is considered optimal because:

- All stream sources are identical, thus for any given sequence number the payload should also be identical.
- Most bit errors should be detected by the CRC-32 algorithm applied to Ethernet, SDH, ATM, etc. These devices will typically discard frames where bit errors occur with the net result being the video ISA will receive a bit error-free stream (though packet loss may/ does occur).
- UDP checksum in verified by the video ISA (after input VQM) and any failures result in a silent discard of the packet.

Initial Sequence Identification

When a service is defined and is enabled (**no shutdown**), the video ISA will monitor for valid RTP packets and on first receipt of a valid RTP packet learn the following information:

- SSRC
- Sequence number
- Timestamp (as timestamp is profile-specific, MPEG2-TS are assumed)

The packet will be inserted into the video ISA playout buffer associated with that particular service and playout when directed (playout algorithm).

Packet Selection

For each valid RTP packet received for a given service will be inserted into the buffer if there is no existing RTP packet that matches the sequence number. Because sequence number and timestamp discontinuities may occur the video ISA makes a limited attempt at validating either as they are not required for MPEG. The video ISA code adopts a philosophy that for the most part sequence number and timestamp increment correctly, but should they prove to be non-contiguous, the packet selection algorithm adapts.

Duplicate packets are detected by sequence number (or timestamp unless M-bit reset it), so should a packet already exist in the buffer with the same sequence number as one received (or one recently played out) it will be discarded. For the purpose of determining recent playout if an incoming sequence number is within 6.25% (- 4096) the packet is considered late and is discarded.

In a multi programme transport stream (MPTS) timestamp is set uniquely for every RTP packet, this is because any RTP packet many contain a number of multiplexed elementary streams. As a result playout is based on the embedded timestamp in each RTP packet. In a single programme transport stream the inverse occurs, many RTP packets can share the same timestamp as it is referenced from the start of picture (and a picture can span many RTP packets). As a SPTS does not contain audio its application is limited to content production and so only MPTS are supported.

Timestamp discontinuities do occur and are normally represented with the Marker bit (M) being set.

Playout time is determined by an internal playout timestamp. The playout timestamp is set independently from the actual timestamp in the packet. The recovered clock is used to determine expected timestamp for every incoming RTP packet.

When a packet is received it is first compared to existing packets in the buffer based on sequence number (assuming here that a stream may be delay hundreds of milliseconds by a backup path yet still be valid); only if this packet is determined to be new RTP packet eligible for buffer insertion will jitter tolerance be evaluated. If jitter tolerance is exceeded then a timestamp discontinuity is

assumed and instead of setting playout timestamp based on the contained RTP timestamp, the actual received time (offset by playout-buffer) is set for the RTP packet playout timestamp.

In normal operation clock is recovered from the timestamp field in the RTP header, is offset by the playout buffer configuration parameter and used to schedule playout of the packet. The playout clock is synchronized with the sender by using an adaptive clock recovery algorithm to correct for wander.

Algorithm summary

- Is the service marked LoT — If a loss of transport occurred, purge the buffer and reset all counters/timers.
- If the service is UP, check the RTP packet sequence number. Compare to sequence numbers contained in the buffer. If no match then check last played sequence number. If the sequence number of this packet is between last played and last played + 4096 then consider this packet late and discard.
- Check the expected timestamp recovered clock value and compare to RTP timestamp: If the expected timestamp is (-ve)jitter tolerance<timestamp<(+ve)jitter tolerance then the packet is admitted to the buffer with a playout timestamp per the embedded RTP timestamp. If jitter tolerance is not maintained this marks a discontinuity event. Set playout timestamp to current clock + playout buffer and enqueue.

Clock Recovery

RFC 2250, *RTP Payload Format for MPEG1/MPEG2 Video*, defines the timestamp format for MPEG2 video streams (which may carry H.264 video): a 90kHz clock referenced to the PCR. Each ingest RTP packet has it's timestamp inspected and it is used in an adaptive clock recovery algorithm. Importantly, these adjustments occur on ingress (not on playout). This serves as a long-term, stable, ingress stream recovered clock.

The 90kHz ingress stream recovered clock is adjusted for each service to account for the encoder's reference clock/difference between the clock in the 7750 SR. This input timestamp is derived from the same RTP packet that is inserted into the buffer, and thus may be subjected to significant jitter. The clock adjustment algorithm must only adjust clock in extremely small increments (in the order of microseconds) over a very long sample period (not bitrate) of at least 30 minutes.

Playout

Playout is the process of regenerating the stream based on playout timestamp.

For each service the operator defines a fixed playout buffer. This serves as an exact offset to the ingress stream recovered clock and serves as playout time for the video ISA. Because timestamp is used for buffer playout, CBR, capped VBR and VBR streams are all supported without pre-configuration. The playout buffer mechanism effectively removes network-induced jitter and restores the output to the rate of the original encoder.

Loss of Transport

In the circumstance that the playout buffer is emptied an LoT is indicated. The video ISA will reset playout timestamp, clock, sequence number, etc., on this event and await the next valid RTP packet for this service.

Video Quality Monitoring

The following terminology is used in this section:

- TNC: Technically non-conformant
- QoS: Quality of Service
- POA: Program Off Air
- Impairment event — A trap/alarm that an impairment event is detected and is termed as tnc. An impaired event is said to have occurred if:
  - → Continuity counter errors were detected.
  - → If PAT /PMT/PCR pids were not present in the video stream for a time period equal to or greater than the configured tnc value in the respective alarm.

    The default value of the impaired threshold in terms of milli second is:

    PAT : 100ms

    PCR : 100ms

    PMT : 400ms
  - → If unreferenced PID is seen in the video stream which has not been referred in the PMT table.
- Impaired seconds — The number of seconds an impaired event was detected.
- Degraded event — A trap/alarm that a degraded event is detected and is termed as QoS. A degarded event is said to have occurred if :
  - → PAT/PMT Syntax error occurs in that second.
  - → Absence of PAT/PMT/PCR pids in the video stream for a time period equal to or greater than the configured qos value in their respective alarms.
    The default value of the degraded threshold in terms of milli second is:

    PAT : 200ms

    PCR : 200ms

    PMT : 800ms
- Degraded seconds — The number of seconds an degraded event was detected.
- Error event — A trap/alarm that an error event is detected and is termed as POA.

- An errored event has occurred if:
  - → If sync loss error has occurred for that particular second. A sync loss is said to have occurred if there are more than 1 consecutive sync byte errors are seen in the stream.
  - → Absence of PAT/PMT/PCR PIDs in the video stream for a time period equal to or greater than the configured poa value.

  The default value of the degraded threshold in terms of milli second is:

  PAT : 500ms

  PCR : 500ms

  PMT : 2000ms
  - → Traffic loss has occurred for that particular second.
  - → Transport error indicator or TEI indicator is set in the transport stream packet header for that particular second in the video stream.
- Errored seconds — The number of seconds an errored event was detected.
- Good seconds — The number of seconds where we do not see any impaired, degraded or errored events.

Pid Stats :

- PID: Displays the value of the pid.
- Is PCR PID : Takes a value Yes/No. If yes, then it indicates that the pid is the PCR PID.
- TEI Err Sec : Counts the number of seconds TEI was set for that particular PID.
- Absent Err Secs: The number of seconds for which the PID was not seen for a particular interval of time which is decided by the alarms set for the Non-Vid Pid Absent and Video PID Absent.
- PID bitrate: Is calculated by counting the number of times the pid occurred in the last second x 188 x 8.

  188 = TS packet size

  8 = Number of bits in a byte
- CC Err Secs: Number of seconds continuity counter errors were seen for that particular PID in the stream.
- PID Type: Specifies that the PID is either video, audio, PAT, PMT, or PCR.
- MPEG Stream Type: If the PID is of video or audio this field informs us about the way the video and audio is encoded.

  For example:
  - → For video : H.264 or Mpeg2 (Only the decimal equivalent defined by the MPEG standard is displayed and not the string)
  - → For Audio : AC-3 or Mpeg-2 (Only the decimal equivalent defined by the MPEG standard is displayed and not the string)

Interval Stats

- Except the PID stats all other stats explained above have interval stats. Information can be obtained about stream status was in the last 1 minute, 5 minute and 15 minute.

MDI - Media Delivery Index (RFC 4445, *A Proposed Media Delivery Index (MDI)*)

- Delay Factor (RFC 4445) — The delay factor is a value which indicates the minimum amount of time a STB buffers to resolve network jitter (i.e., it is the minimum STB buffer depth in ms). RTP timestamp will be used as the definitive time indicator (the notional drain rate).
- Loss Rate (RFC 4445) — The Media Loss Rate is the number of media (Transport Stream) packets lost over a certain time interval. This is reported in TS/sec. Each RTP packet lost is assumed to have 7 TS packets lost.
- In absence of traffic MDI values will be reported as N/A . These stats are reported over current (current second) , 1 minute, 5 minutes and 15 minutes intervals

In many instances IPTV operators are unable to identify the cause of visual impairments which are present in almost every video distribution network because the IPTV network has so many moving parts While head end transport-stream monitoring; full reference video analysis (comparing the source content to the encoded output), and; STB probes allow an operator to establish whether the contribution source, the encoder, or the network is the problem the network is a very complex thing.

Operators can use another measurement point in the network, just prior to the last mile such that network faults can be characterized as being between the head end and last mile (transport) or in the last-mile itself.

The multicast video quality monitoring solution provides an inspection point for the multicast video stream that is combined with other analysis methods to create a full view of video issues and help troubleshoot the part of the network causing the issue.

Video quality monitoring is one part of a video assurance program and is combined with:

- TS analysis on the encoder output (to detect encoder errors);
- Full-reference PSNR and PQR on the encoder output (to detect over-encoding, noise and other contribution or encoding artefacts)
- STB reporting (such as packet-loss, RET events, packet errors) from the entire STB population
- STB probes performing full-reference monitoring (against test streams)
- STB probes performing channel-change times, estimated PSNR, etc

Multicast video monitoring within the network can be positioned as complementary to STB reporting and head end analysis, and but should not attempt to perform either of these functions.

Because the network node is not capable of decrypting a MPEG transport stream is primarily used to identify correctable and un-correctable network errors, correlate them with network events (i.e., routing reconvergence, interface failure, etc) and provide summary reports and alarms.

For operators who do not have existing STB probes or reporting, a network-based VQM solution can provide insight into quality issues the network may be contributing to, possibly reducing the amount of STB probe investment that is needed. (i.e., both probes and the 7750 VQM reports many of the same issues in terms of picture quality, fewer probes are needed to test channel change delay, etc).

The metrics which VQM can report are based on the use of RTP streams which provide per-packet sequencing and an indication of picture type. These two parameters along with measured bitrate allow VQM to produce estimated MOSv scores for both stream ingress (uncorrected) and stream egress (corrected) outputs.

Reportable metrics include:

- Relevant SCTE-143 error counters
  - → PAT
  - → PMT
  - → PCR
  - → Transport errors, etc
- ETSI TR 101 290
  - → PID
  - → SI repetition
  - → Degraded blocks/intervals, etc
- MDI (RFC 4445)
- RTP Measurements (RFC 3357, *One-way Loss Pattern Sample Metrics*)
- Forwarded and impaired I-/B-/P-frame counts
- GOP length
- Video/audio/stream bitrate

These metrics are collected per stream and have relevant parameters (such as profile and PIDs) pre-defined, these will be collected into a so-called stream ID. Reports (containing numeric metrics) and alarms (log, SNMP or syslog) can be generated.

For each group, reports contain:

- Stream ID (S,G / SSRC)
  - → Stream A (ingress)
    - – Statistics
  - → Stream B (ingress)
    - – Statistics
  - → Output
    - – Statistics

Reports are non-realtime and are compiled into an XML format for FTP extraction with a resolution of less than 5 minutes.

Event alarms are reported by log, syslog or SNMP (existing log interface).

VQM is an optional module available on the input side, or output side of the video ISA. On input, it is applied prior to ad-insertion, H-RET, and duplicate stream protection, conversely when on the output side it is applied only to multicast streams after ad-insertion, H-RET and duplicate stream protection.

Because of the large number of channels and the nature of measuring input and output sides, VQM is highly reliant on the use of RTP extensions to provide relevant transmission metrics to the VQM analysis module. In a typical head end a multicast stream will be scrambled to encrypt its video and/or audio. When this encryption occurs, it is typical for the entire payload of the transport stream (for the nominated PID) to be completely scrambled. The consequence of such is that the video and audio PES headers, which reveal much about the picture and timing information, are unavailable to the VQM program.

VQM utilizes intelligent RTP re-wrapping. RTP re-wrapping is a prerequisite for ad insertion and Fast Channel Change (FCC) and involves marking packets before encryption based on the picture type (most importantly, the start of the I frame of IDR frame in H.264.

The Alcatel-Lucent VSA as currently defined, re-multiplexes each transport stream into a new RTP packet. By doing so it allows the separation of different picture types into their own respective RTP packets, and the separation of audio packets from video packets to allow different synchronisation in events of FCC. In effect, it pulls the elementary streams back into their component forms while retaining the syntax and structure of the MPTS.

For information about Alcatel-Lucent VSAs, refer to the 7750 SR OS System Management Guide.

Meanwhile, additional information can be made available, prior to scrambling, of the picture information for quality analysis. The quality analysis performed by the VQM module emphasizes impairments caused by network issues and transport stream syntax given the relative proximity of the router to the customer.

When the video ISA is deployed alongside the ALU VSA re-wrapper a custom RTP header extension is sent with each RTP packet.

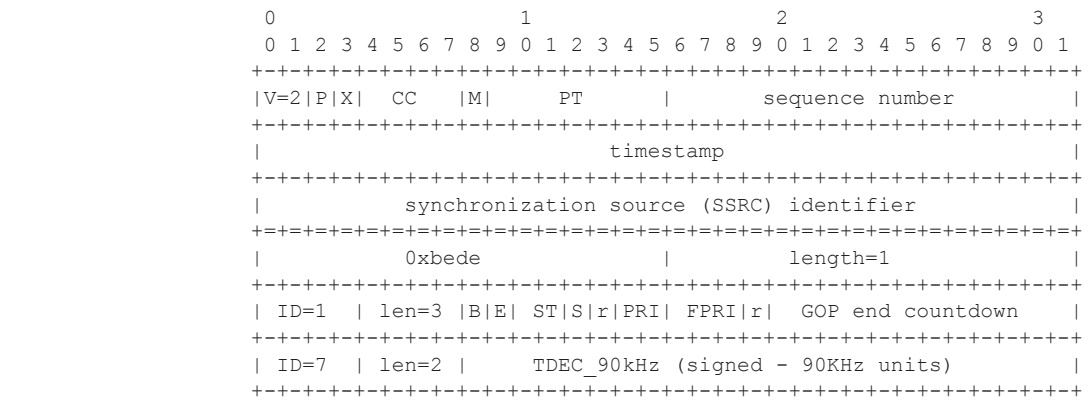

Figure 31: RTP Header Extension

Where:

```
B (Frame Begin Flag): set if a frame starts in this packet
 E (Frame End Flag): set if a frame finishes in this packet
 ST (Stream Type)
 00 video
 01 audio
 10 data/padd/other
 11 Reserved
 S (Switch bit): set to 1 in all RTP packets from the moment
 the client should do the IGMP join (rewrap does not fill it)
 r: reserved (set to 1)
 PRI: Packet Priority (coarse)
 FPRI: Fine-grained priority
 PRI FPRI dec DSCP
 --- ---- --- ----
 3 7 31 AF41 Video IDR frame
 3 0 24 AF41 Audio
 2 0 16 AF41 Reference frame (P in MPEG2, I or P or some Bs in AVC)
 1 7 15 AF42 Non-reference frame (most Bs in MPEG2, some Bs in AVC)
 1 5 13 AF42 Trans-GOP frames, undecodable in some circumstances (some Bs in MPEG2)
 0 4 4 AF43 Rest of cases (data, secondary videos, etc)
 0 1 1 AF43 Padding packets
where AF41=100010, AF42=100100, AF43=100110 (DSCP bits in the IP header)
```

VoIP/Video/Teleconferencing Performance Measurements

The feature provides ability to measure and provide statistics to allow reporting on voice and video quality for VoIP and teleconferencing (A/V) applications. A sampled deployment is shown in the picture below (Figure 32). Although a distributed model is shown, a hub-and spoke model, with AA-ISA deployed only on one side of the traffic flow, is also supported.

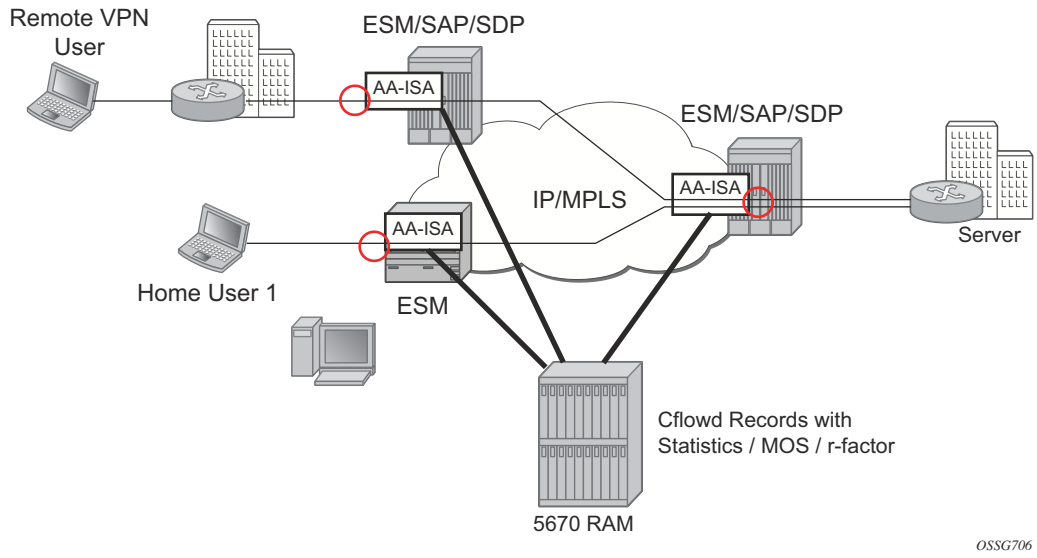

Figure 32: Voice/Video Monitoring Deployment Example

Because of network-based AA, the operator has an ability to monitor voice, video, teleconferencing applications for a given AA subscriber regardless of the type of that subscriber (a residential subscriber vs. a user of a business VPN service). AA-ISA monitors UDP/RTP/RTCP/ SDP headers for each initiated call/application session (sampling may be provided – although, it is expected that a sampling rate will be smaller than that of TCP-applications due to the nature of the voice/video applications – longer lasting and smaller number of sessions/calls per subscriber). AA ISA gathers statistics and computes MOS-scores/R-factor results per each call/ application session. At the end of a call (/application session closure), AA-ISA sends the statistics and computed scores to a Cflowd collector (the Cflowd infrastructure was introduced for TCP-performance but modified to carry voice/video specific data is used). The collector summarizes and presents the results to the operator/end user.

Mean Opinion Score (MOS) Performance Measurements Solution Architecture

AA-ISA integrates a third party MOS software stack to perform VoIP and video MOS measurements. This software provides:

- Call quality analysis using optimized ITU-T G.107
- Measurements of perceptual effects of burst packet loss and recency using ETSI TS 101 329-5 Annex E Extensions
- Measurements and analysis of RTCP XR (RFC3611) VoIP metrics payloads.

AA software monitors the associated SDP channel and passes codec information (when available) to the subsystem which monitors VoIP. The video bearer channels traffic generates a wide variety of A/V performance metrics such as:

- Call quality metrics
  - → Listening and conversational quality MOS scores – MOS-LQ, MOS-CQ
  - → Listening and conversational quality R-factors – R-LQ, R-CQ
  - → Estimated PESQ scores – MOS-PQ
  - → Separate R-factors for burst and gap conditions – R-Burst, R-Gap
  - → Video MOS-V and Audio MOS-A
  - → Video Transmission Quality - VSTQ
- Video stream metrics
  - → Good and impaired I, B, P, SI, SP frame counts
  - → Automatic detection of GoP structure and other key video stream attributes such as image size, bit rate, codec type
- Transport (IP/RTP) metrics
  - → Packet loss rate, packet discard rate, burst/gap loss rates
  - → Packet delay variation/ jitter
- Degradation factors
  - → degradation due to loss, jitter, codec, delay, signal level, noise level, echo, recency

Once a flow terminates, AA software retrieve the flow MOS parameters from the subsystems, formats the info into a Cflowd record and forwards the record to a configured Cflowd collector (RAM).

RAM collects Cflowd records, summarizes these records using route of interest information (source/destinations). In addition, RAM provides the user with statistics (min/max/ avg values) for the different performance parameters that are summarized.

Retransmission and Fast Channel Change

RET and FCC Overview

The following sections provide an overview of RET and FCC.

Retransmission

Retransmission (RET) for RTP (RFC 3550, *RTP: A Transport Protocol for Real-Time Applications*) is based on a client/server model where the client sends negative acknowledgments (NACKs) using Real-time Transport Control Protocol (RTCP) (RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*) to a RET server when the client detects missing sequence numbers n the RTP stream. The RET server which caches the RTP stream, for example in a circular buffer, detects missing sequence numbers in the replies to the NACKs by resending the missing RTP packets as illustrated in Figure 33.

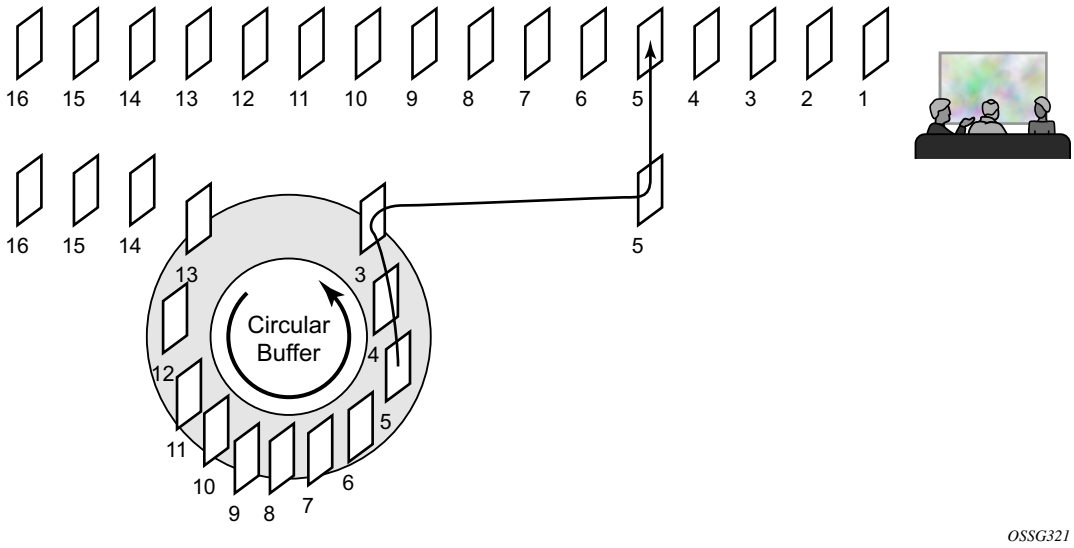

Figure 33: RET Server Retransmission of a Missing Frame

The format of the reply must be agreed upon by the RET client and server and can be an exact copy (Payload Type 33 as defined in RFC 3551, *RTP Profile for Audio and Video Conferences*

with Minimal Control) or sent with a different Payload Type using an encapsulating RET header format (RFC 4588, *RTP Retransmission Payload Format*).

RET has been defined in standards organizations by the IETF in the above-noted RFCs and Digital Video Broadcasting (DVB) in “Digital Video Broadcasting (DVB); Transport of MPEG-2 TS Based DVB Services over IP Based Networks (DVB-IPTV Phase 1.4)” which describes the STB standards.

STBs that have a port of the Alcatel-Lucent RET/FCC Client SDK are an example of a standards-compliant RET Client implementation.

Fast Channel Change (FCC)

FCC is an Alcatel-Lucent method based on a client/server model for providing fast channel changes on multicast IPTV networks distributed over RTP. During a fast channel change, the FCC client initiates a unicast FCC session with the FCC server where the FCC server caches the video stream and sends the channel stream to the FCC client starting at the beginning of a Group of Pictures (GOP). Beginning at a GOP in the past minimizes the visual channel transition on the client/STB, but the FCC unicast stream must be sent at an accelerated rate in the time domain to allow the unicast stream to catch up to the main multicast stream, at which point, the FCC server signals to the client to join the main RTP stream.

Figure 34 illustrates the FCC client and server communication.

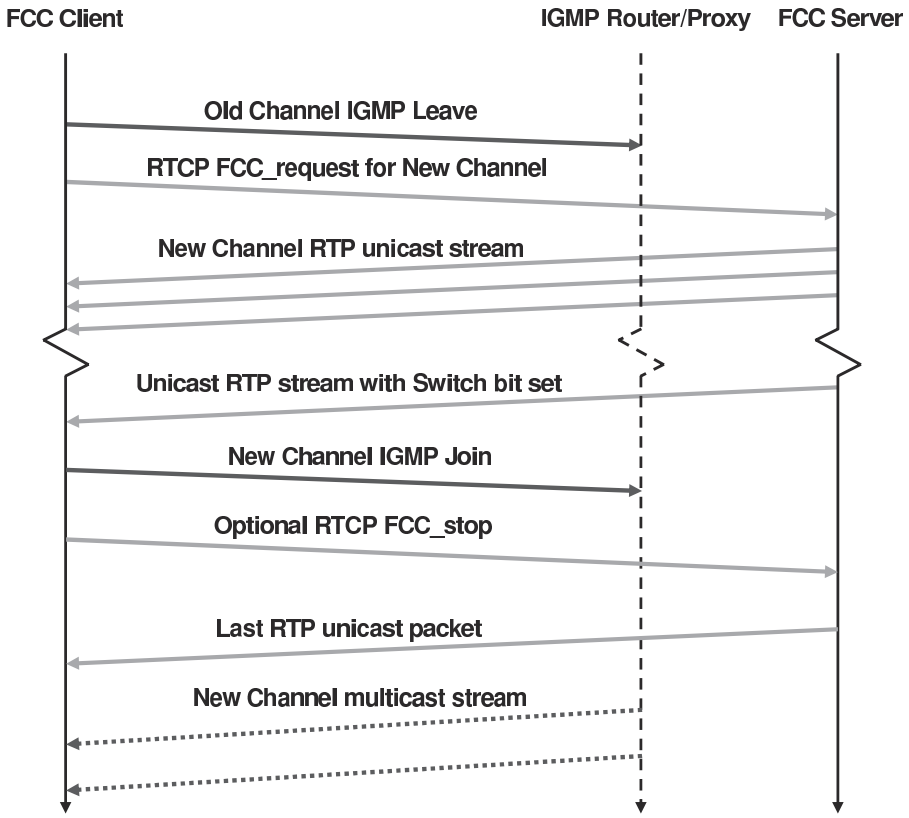

Figure 34: FCC Client/Server Protocol

There are two techniques for compressing the FCC unicast stream in time to allow the unicast session to catch up to the multicast stream: bursting and denting. When bursting, the stream is sent at a rate faster than multicast stream, for example, the stream can be "bursted" at 130% (or 30% over the nominal) multicast rate. "Denting" is a technique where less important video frames are dropped by the FCC server and not sent to the FCC client. Hybrid mode combines bursting and denting.

Bursting is illustrated in Figure 35 and denting is illustrated in Figure 36.

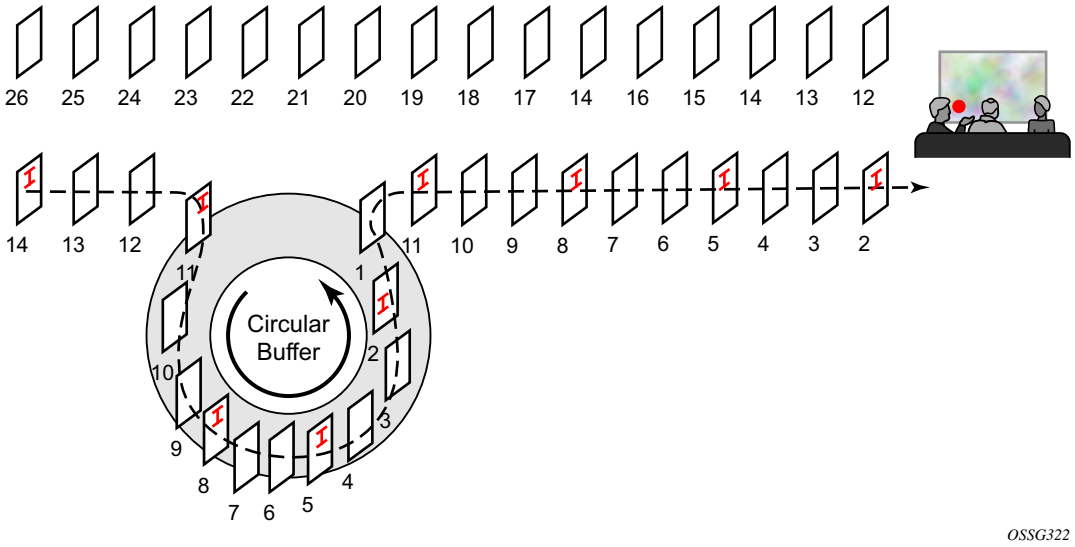

Figure 35: FCC Bursting Sent Faster Than Nominal Rate

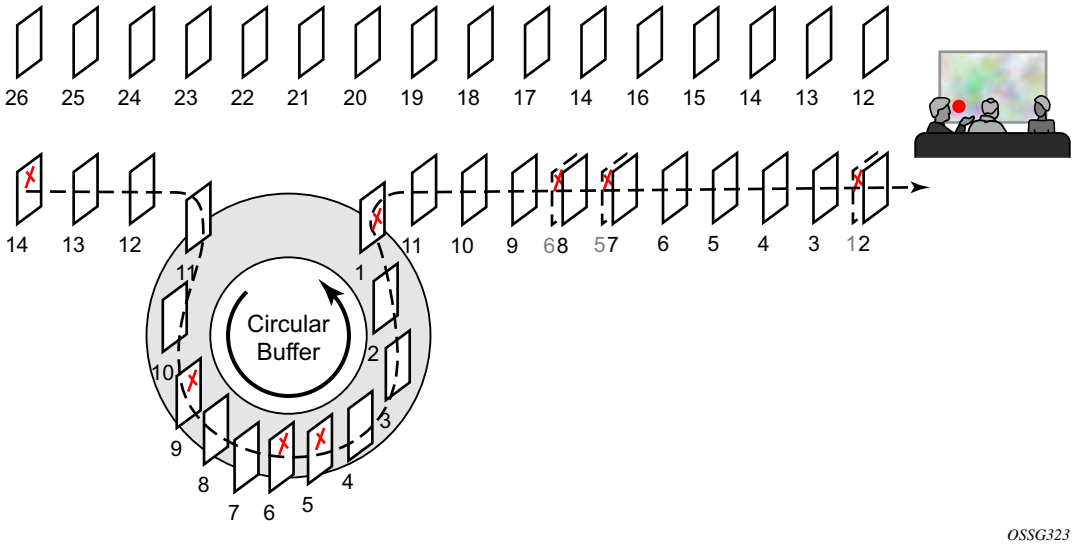

Figure 36: FCC Denting Removing Less Important Frames

When the unicast session has caught up to the multicast session, the FCC server signals to the FCC client to join the main multicast stream. The FCC server will then send the unicast session at a lower rate called the “handover” rate until the unicast session is terminated.

Note that the FCC server functionality requires the Alcatel-Lucent 5910 Video Services Appliance (VSA) Re-Wrapper which is used to encapsulate and condition the multicast channel streams into RTP, adding important information in the RTP extension header. Also, the ISA FCC server requires an STB FCC client based on the Alcatel-Lucent FCC/RET Client SDK.

Retransmission Client

The ISA RET client is used in hierarchical RET deployments and performs upstream corrections for missing packets in the RTP multicast stream to ensure that the RET server has all the packets for the stream.

The RET client is supported within a VPLS, IES or VPRN service context as applicable to the platform. The RET client source address is explicitly assigned. In a VPLS, the RET client IP appears to be an IP host within the service, and like a host, the RET client is also configured with a gateway IP address to provide a default route to reach the upstream RET server.

Whenever the RET client receives a retransmission from an upstream RET server, the replies are sent downstream as multicast in the multicast service using Payload Type 33 which is the Payload Type for an original stream.

Whether the RET client is active for a given multicast channel is defined in the multicast information policy where channels are defined. The channel configuration for the RET client within the policy is an explicit enable/disable of the RET client and the IP address and UDP port for the upstream RET server for the channel.

The ISA RET server supports the network model where there are separate service instances for unicast and multicast traffic that are cross-connected and multicast replicated downstream in the network, for example, where an access node provides the multicast service cross connect and replication at the last mile. If there are separate multicast and unicast service instances, the multicast service instance must be configured in the unicast service, and the unicast and multicast services must use the same multicast information policy.

Retransmission Server

The ISA RET server is supported within a VPLS, IES or VPRN service context as applicable to the platform.

Whether the RET server is active for a given multicast channel is defined in the multicast information policy where channels are defined. The channel configuration for the RET server within the policy is an explicit enable/disable of the local RET server (that is, whether the channel should be buffered), the RET buffer size for the channel in the ISA and a channel type (Picture-in-Picture (PIP), Standard Definition (SD) or High Definition (HD)). The RET buffer should be large enough to account for the round trip delay in the network; typically, a few hundred milliseconds is sufficient.

In a VPLS service, a single IP address is assigned to the RET server, and it acts like an IP host within the service.

In an IES or VPRN service, up to 16 IP addresses can be assigned to a video interface.

The video policy within the multicast information policy defines the characteristics for the how the RET server should respond to NACKs received on an IP address. The different characteristics defined in a RET server "profile" are for each channel type (PIP, SD and HD):

- Enable/disable for the RET server (that is, whether requests should be serviced or dropped).
- The RET rate (as a percentage of the nominal channel rate).

Typically, RET replies are sent below line rate because most dropped packets occur in the last mile and sending RET replies at a high rate may compound any last mile drop issues.

The IP address(es) of the RET server is(are) defined in the unicast service instance, whereas the UDP port for the RET server is defined in the "default" bundle in the multicast information policy. The same UDP port is used for all RET server IP addresses that use the particular multicast information policy.

The ISA RET server supports the network model where there are separate service instances for unicast and multicast traffic that are cross-connected and multicast replicated downstream in the network. If there are separate multicast and unicast service instances, the unicast and multicast services must use the same multicast information policy.

Fast Channel Change Server

The ISA FCC server is supported within a VPLS, IES or VPRN service context as applicable to the platform. VPRN services are not supported on the 7450 ESS.

Whether the FCC server is active for a given multicast channel is defined in the multicast information policy where channels are defined. The channel configuration for the FCC server within the policy is an explicit enable/disable of the local FCC server (that is, whether the channel should be buffered) and a channel type PIP, SD or HD. When FCC is enabled, three (3) GOPs are stored in the buffer. the channel also defines an optional fcc tuning parameter called the fcc Minimum Duration which is used by the FCC server to determine which GOP to start the FCC unicast session. If there are too few frames of the current GOP stored in the fcc server buffer (based on number of milliseconds of buffering), the FCC server will start the FCC session from the previous GOP.

In a VPLS service, a single IP address is assigned to the FCC server, and it acts like a IP host within the service.

In an IES or VPRN service, up to 16 IP addresses can be assigned to a video interface.

The Video Policy within the multicast information policy defines the characteristics for the how the FCC server should respond to FCC requests received on an IP address. The different characteristics defined in an FCC server "profile" are for each channel type (PIP, SD and HD):

- Enable/disable for the FCC server (for example, should the requests be serviced or dropped).
- The FCC mode: burst, dent or hybrid.
- The burst rate (as a percentage above the nominal channel rate) for PIP, SD and HD channel types.
- The multicast handover rate (as a percentage of the nominal channel rate) used by the server after it has signaled the client to join the main multicast channel.

Different FCC rates are allowed for each of the channel types because the channel types have different nominal bandwidths. For example, the last mile may only be able to reliably send a 25% burst (above nominal) for HD whereas the equivalent bit rate for SD is a 75% burst. The profiles are designed to provide flexibility.

The IP address of the FCC server is defined in the unicast service instance, whereas the UDP port for the FCC server is defined in the "default" bundle in the multicast information policy. The same UDP port is used for all FCC server IP addresses that use the particular multicast information policy.

The ISA FCC server supports the network model where there are separate service instances for unicast and multicast traffic that are cross-connected and multicast replicated downstream in the

network. If there are separate multicast and unicast service instances, the unicast and multicast services must use the same multicast information policy.

Logging and Accounting for RET and FCC

In previous releases, logging and statistics were maintained for active sessions (RET and FCC).

This feature now provides more permanent logging, statistics and accounting for:

- RET Server sessions stats
- FCC session stats
- ADI events

RET Server Session Stats

For RET Server Stats, the RET session table entries will be sampled and periodically written to XML accounting records.

The basic framework is (requiring a CLI and perhaps some additional tuning) is:

- Session statistics will be written to a record in an XML file on a periodic basis with the sample period being 5 minutes or longer.
- Session statistics are written to a record when a) the session is removed from the session table, b) if the session is exists for more than two write periods.
- All statistics will be the total values (that is, not incremental values across sampling periods).

RET and FCC Server Concurrency

Even though the previous sections discussed the RET server and FCC server as separate entities, the ISA can support RET and FCC servers at the same service at the same time. As such, the configuration commands and operational commands for the services are intermingled. If both the RET server and FCC server are enabled for a given channel, a single buffer is used for caching of the channel.

A maximum bandwidth limit for all server requests can be defined for a given "subscriber" which is equated with the source IP address. Before an ISA server processes a request, the ISA calculates the bandwidth to the subscriber required, and will drop the request if the subscriber bandwidth limit will be exceeded.

The ISA services RET and FCC requests on a first in, first out (FIFO) basis. Before servicing any request, the ISA calculates whether its egress bandwidth can handle the request. If there is insufficient egress bandwidth to handle the service request, the request is dropped. Near the ISA's egress limits, RET requests will generally continue to be serviced whereas FCC requests will be dropped because RET sessions are generally a fairly small percentage of the nominal rate and FCC sessions are slightly below to above the nominal channel rate.

Prerequisites and Restrictions

This section summarizes some key prerequisites and restrictions for the RET client, RET server and FCC server.

- Both RET and FCC require RTP as the transport stream protocol.
- FCC requires the Alcatel-Lucent 5910 VSA Re-Wrapper.
- FCC requires an implementation of the Alcatel-Lucent 5910 STB Client.
- The multicast information policies must be the same on multicast and unicast services which are cross connected downstream.
- Support for up to four ISA-MSs in a video group
- Only a single IP address and profile are supported within a VPLS service for RET or FCC, so only a single Profile can be supported in a VPLS service.
- Up to 16 IP addresses can be configured for a Layer 3 service video interface (IES or VPRN) with each supporting a distinct profile.
- There can be a maximum of 32 IP addresses across all Layer 3 service video interfaces per chassis.

Multi-Service ISA Support in the IOM-3 for Video Services

In previous releases, the Multi-Service ISA was supported in the iom-20g-b and the iom2-20g for video services. Now, this feature provides support for the Multi-Service ISA when installed in an iom3-xp card on both the 7450 ESS and the 7750 SR.

Prioritization Mechanism for RET vs. FCC

In previous releases, RET and FCC requests are processed with the same priority. Since RET generally has a more direct impact on a subscriber's "quality of experience", service providers are prioritizing RET as a feature over FCC, and for those that want to implement both, the preference is to have a mechanism to prioritize RET over FCC when there is contention for resources.

Now, this feature provides a mechanism to reserve an explicit amount of egress bandwidth for RET for all the ISAs within an video group. If the amount of egress bandwidth is less than the reserved amount, FCC requests are discarded and only RET requests processed. The bandwidth will need to be dynamically adjusted per ISA within the video group if ISAs become operational/non-operational within the group.

RET Features

Statistics ALU SQM MIB Additions

Alcatel-Lucent in Portugal has developed a network management application that does a statistical analysis of retransmissions to analyze the video quality. The following are existing MIB entries.

- TmnxVdoSessionEntry ::= SEQUENCE {
- tmnxVdoSessionSourceAddrType InetAddressType,
- tmnxVdoSessionSourceAddr InetAddress
- tmnxVdoSessionSourcePort InetPortNumber,
- tmnxVdoSessionSSRCId Counter32,
- tmnxVdoSessionUpTime Unsigned32,
- tmnxVdoSessionExpireTime Unsigned32,
- tmnxVdoSessionCName TNamedItem,
- tmnxVdoSessionDestAddrType InetAddressType,
- tmnxVdoSessionDestAddr InetAddress,
- tmnxVdoSessionRxFCCRequests Counter32,
- tmnxVdoSessionTxFCCReplies Counter32,
- tmnxVdoSessionTxFCCPackets Counter32,
- tmnxVdoSessionTxFCCOctets Counter32,
- tmnxVdoSessionRxRTRequests Counter32,
- tmnxVdoSessionTxRTReplies Counter32,
- tmnxVdoSessionTxRTPackets Counter32,
- tmnxVdoSessionTxRTOctets Counter32

The following are new entries:

- Total number of sequences of 10 — total sequences of 2 to 10 lost packets
- Total number of sequences of 20 — total sequences of 11 to 20 lost packets
- Total number of sequences of 30 — total sequences of 21 to 30 lost packets
- Total number of sequences of 40 — total sequences of 31 to 40 lost packets
- Total number of sequences of more ?total sequences of 41 or more lost packets}

RET Server Multicast Tuning Parameters

Downstream RET requests are responded to using multicast when there are a number of identical RET requests with the assumption that there was a loss in the network that affected a number of clients. In this instance, the retransmitted frames will be sent as Payload Type 33 as original packets and not in the RFC 4588, *RTP Retransmission Payload Format*, retransmission format.

The **rt-mcast-reply** command can tune the RET server as to when to use multicast to reply to RET requests have the option to disable multicast responses.

FCC Features

FCC Hybrid Mode Support

There are three modes of operation supported for FCC:

- In burst mode, the unicast FCC traffic is sent faster than nominal rate (bursted above nominal).
- In dent mode, packets are dropped from the unicast FCC stream based on a defined threshold for markings added to the packet that indicate the importance of the packet to the audio/video stream added by the rewrapper.
- Hybrid mode combines both bursting and denting.

Ad Insertion

Local/Zoned Ad Insertion

Transport Stream Ad Splicing

Alcatel-Lucent's Local/Zoned ADI feature allows a 7750 SR with the ISA-MS (the "splicer") to perform ad splicing in an MSTV environment. The splicer is a post-A server transport stream (TS) splicer and can splice into encrypted or unencrypted transport streams. The splicer is positioned between the A-server and the D-server. Figure 37 shows an ad insertion model displaying components.

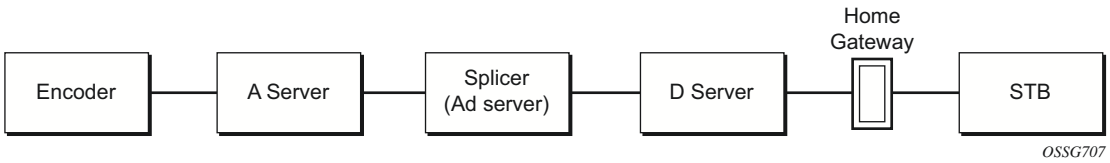

Figure 37: Ad Insertion Model

The ad insertion process is initiated when the splicer detects the SCTE 35 cue signal that identifies the upcoming start and end of the advertising time slot. The splicer communicates with the ad server using SCTE 30 standard messaging and will be instructed by the ad server:

- To take advantage of an ad insertion opportunity or avail and
- Determine the ad to be spliced into the main stream, if applicable.

The ad servers must be configured for ad content to match encoder configurations for video/audio streams. The ad server sends the ad stream to the ad splicer and the ad splicer will switch it into the main stream as dictated by the digital splice points (Figure 38). The ad splicer can splice multiple ads into multiple channels simultaneously.

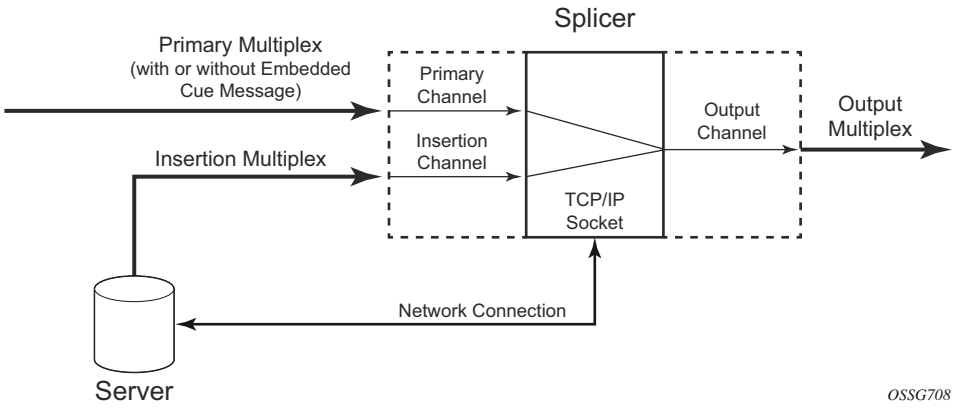

Figure 38: Transport Stream Ad Splicing

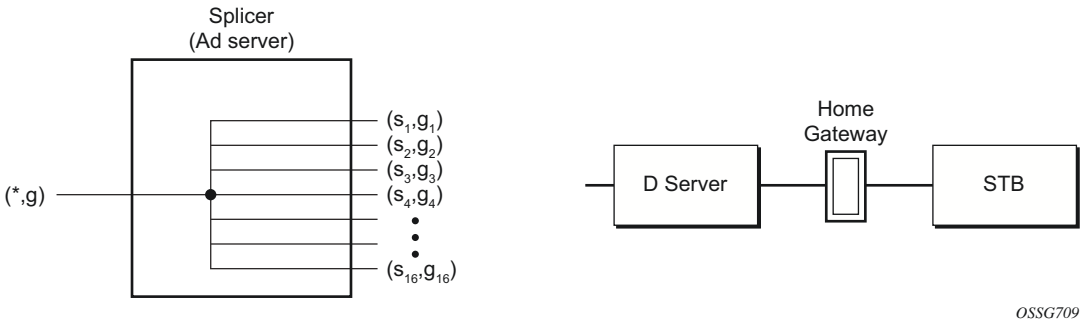

Figure 39: Splicer Model

Note that IPTV encryption and Digital Rights Management (DRM) can be applied to the transport stream payload but not to the transport stream (TS) header which allows a TS splicer to splice into encrypted streams, although the spliced ad content will in all cases be unencrypted. TS splicing does not put any requirements on the middleware platform as ad insertion will be outside the middleware's knowledge and control.

The Figure 40 depicts a TS flow with various MUXed elementary streams (ES) identified by a unique Packet Identifier (PID). The Program Map Table (PMT) is used as the legend to map PID to elementary streams. The digital cue points are also identified by separate unique PID also defined in the PMT that is used by the TS splicer to know when to splice-in and splice-out of the stream. It is important to note that the only important thing that a TS splicer needs are the headers of the TS packets, and the underlying payload of each ES is not needed. This gives the splicer flexibility and makes it agnostic to the ES payload types.

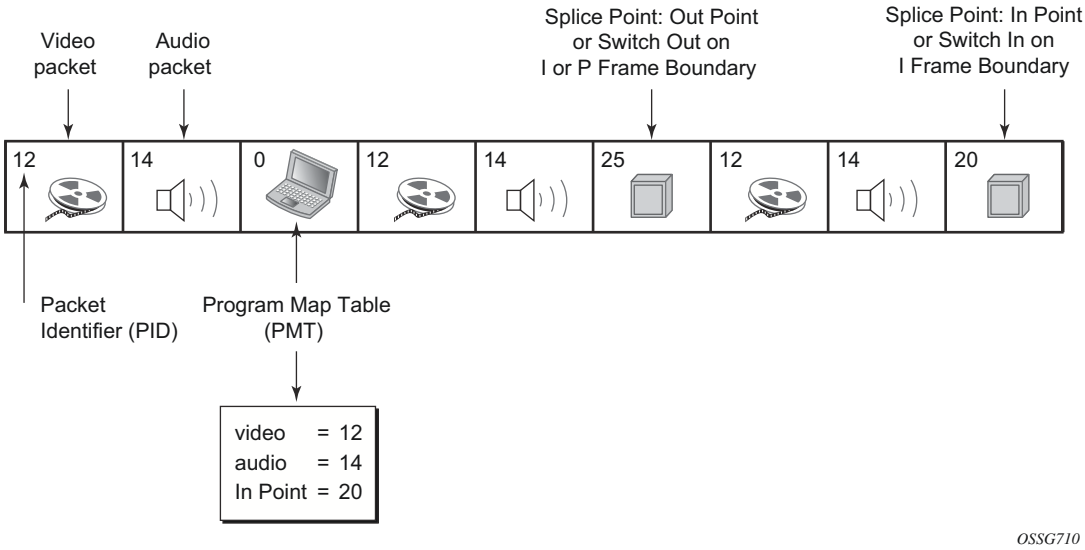

Figure 40: Transport Stream Flow Example

Ad Zones

Within the splicer, zones are created by taking an ingress main channel multicast group, for example (*,G) or (S,G), and creating one or more egress “zone channels” on distinct source-specific multicast (SSM) groups (S1,G1), (S2,G2), etc. Up to 16 zones can be configured for each ingress multicast channel. The group multicast address for the zone channels need not be unique and can actually be the same as the ingress channel, but the SSM sources for the zone channels must be distinct.

Within SCTE 30, the main channel and zone channel are identified by an ASCII string name. These names must be unique and will be used when the splicer communicates with the ad server.

The input stream can be depicted through the following semantics diagram.

CHANNEL1	→	CHANNEL1_North (S1, G1)
(S, G)	→	CHANNEL1_South (S2, G2)
	→	CHANNEL1_East (S3, G3)
	→	CHANNEL1_West (S4, G4)
	→	CHANNEL1_Central (S5, G5)

where (S,G) is the input main channel stream mapping into five (5) (Sx, Gx) which are zone channel streams.

S1..S16 must be IP addresses in the video interface subnet but not the video interface address itself. This implies that traffic for the zones will be sourced from the ISA-MS.

To facilitate traffic from (S,G) to go to the ISA-MS, a static IGMP (S,G) must be configured on the video interface.

Local/Zoned ADI Prerequisites and Restrictions

This section describes prerequisites and restrictions for the local/zoned ADI feature:

- Network Time Protocol (NTP) is required to keep time synchronized between the ad server and the splicer. The time synchronization system helps keep the splicer and the server within +/-15 ms of each other.
- ADI is only supported within a Layer 3 IES or VPRN service.
- Splicing an SD advertisement into an HD main stream is supported, but splicing of an HD advertisement into an SD is not supported.
- The SCTE 30 connection between the ad server and the splicer must be maintained on separate IP addresses on the splicer within the video service.
- Up to 2 ad servers can be configured for redundancy.
- ADI only supports a single ISA-MS member in a video group.
- Up to 16 zone channels can be configured for a main channel.
- The audio re-ordering value in the multicast information policy must match the audio re-ordering configured on the A Server for reliable audio splicing.
- For best results, the ad should start/end with few frames of muted audio.
- The frequency of IDR frames in the network and ad streams must be less than one IDR frame every 1.3 seconds.
- Only the **splice_insert** command of SCTE-35 cue message is supported. The **splice_immediate** command is not supported.