Agama

WHITE PAPER

AGAMA ANALYZER

Seeing through SRT

SRT repairs packet loss and encrypts live video. Both features make monitoring harder. This paper explains how the Agama Analyzer monitors SRT links, including passive monitoring of an existing session.

Agama

Executive summary

SRT carries live video between pipeline components when managed multicast is not available. It encrypts the payload and repairs packet loss within a set latency window. These features make monitoring harder. Recovery can hide packet loss until the window runs out and the picture breaks. Probes cannot inspect the encrypted video. Each endpoint sees only its side of the link. A separate monitoring receiver sees a different session.

The Agama Analyzer supports two ways to monitor SRT links. It can act as an endpoint in a monitoring session or passively capture traffic from a mirror port. The passive tap captures both directions of the actual session. It uses the session passphrase to decrypt the payload and reconstruct the stream played by the receiver. It measures the time left for packet recovery, where loss occurs and the health of the return path. The Analyzer then runs its full analysis on the reconstructed stream, using the same views and platform as the rest of the video chain.

Where SRT fits

Within a facility, video often travels over a managed multicast network. Links between facilities, cloud services and remote productions face less predictable conditions. Packet loss and jitter can disrupt live video on these paths.

SRT handles these conditions across three common uses:

Contribution	An encoder at an event sends video to a facility or cloud service.
Delivery between components	SRT connects gateways, transcoders, sites and cloud services.
Distribution	A listener serves partners or affiliates, with a separate session for each receiver.

The operator is responsible for the service but may control only one endpoint. The network between the endpoints may belong to someone else.

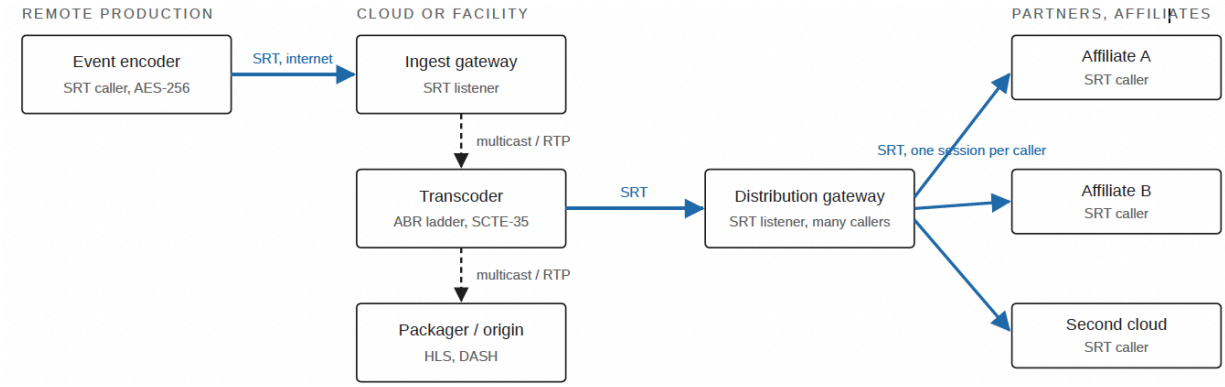

Figure 1. A live video pipeline. In this example, SRT connects the remote production, facility and distribution partners. Internal links use managed multicast.

How SRT works

In a typical live video session, media travels over UDP in one direction. The receiver sends feedback in the other.

Each data packet carries a sequence number and a sender timestamp. When the receiver detects a missing packet, it reports the loss. The sender then retransmits the packet.

The receiver schedules playback using the sender timestamp and a negotiated latency. This delay gives SRT time to report loss and recover missing packets. If a packet misses its playback deadline, the receiver skips it.

For encrypted sessions, the passphrase enables the endpoints to establish the keys used to encrypt the payload. Control packets remain readable. A monitor can therefore inspect protocol activity without decrypting the video. It needs the passphrase to inspect the content.

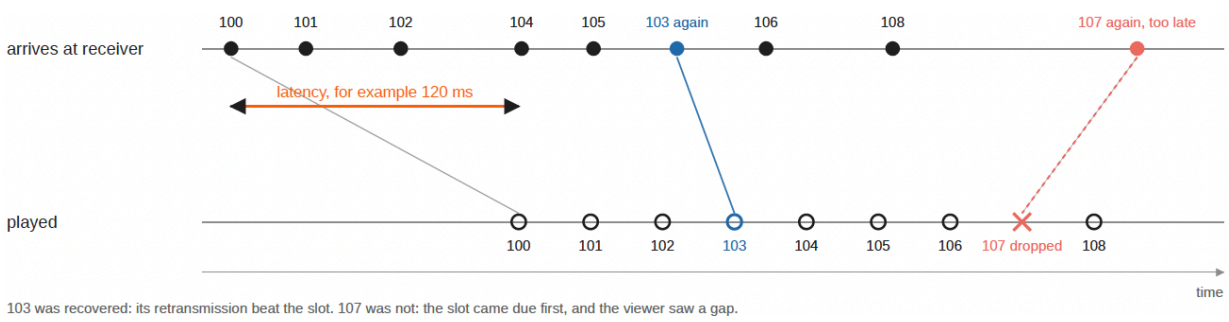

Figure 2. Packet recovery must finish before the playback deadline. Packet 103 arrives in time after retransmission. Packet 107 arrives too late and is dropped. The time remaining before the deadline is the latency headroom.

Agama

Why SRT is hard to monitor

Recovery can hide deteriorating conditions

On a multicast link, packet loss causes immediate continuity errors. SRT can recover that loss before it affects the output.

The picture may stay clean while retransmissions increase and latency headroom falls. Once recovery takes too long, the receiver drops packets. Monitoring only the output can miss the warning signs.

Endpoint statistics show part of the picture

Each endpoint reports its own packet counts, retransmissions, drops and round-trip estimates. Vendors may expose these values in different formats.

The sender does not directly observe playback. The receiver detects missing packets but cannot locate the loss along the network path.

Encryption limits conventional probes

A conventional probe on a mirror port sees UDP traffic but cannot inspect the encrypted payload. It cannot check transport-stream continuity, decode the video or assess the picture.

A monitoring receiver sees a separate session

Some listeners accept only one caller. Others allow an additional monitoring receiver, but that receiver adds traffic and creates a separate session.

The monitoring session has its own packet loss and timing. Its results may differ from those of the service being monitored.

Operators therefore need to answer two questions:

1. What did the receiver play after SRT recovered or dropped packets?
2. Did loss occur before or after the observation point, or on the return path carrying receiver feedback?

QUESTION	SENDER STATISTICS	RECEIVER STATISTICS	CONVENTIONAL NETWORK PROBE	AGAMA SRT TAP
Is the session losing packets?	Based on receiver reports	Shows incoming loss	Limited without SRT analysis	Observes media and feedback
How much latency headroom remains?	Limited visibility	Depends on available metrics	Not available	Measures headroom
Is loss before or after the observation point?	Cannot locate it	Cannot locate it	Not available	Distinguishes loss before and after the tap
Is the return path healthy?	Partial visibility	Limited visibility	Limited without SRT analysis	Analyses feedback-path health
What stream reached playback?	Cannot observe it directly	Statistics describe the outcome	Cannot inspect encrypted content	Reconstructs and analyses the played stream

How the Agama Analyzer monitors SRT

Monitoring as an endpoint

The Analyzer can act as an SRT caller or listener. Use this mode when a monitoring session is available—for example, a dedicated monitoring feed, a listener with spare capacity or a test point at a service handover.

Monitoring through a passive tap

SRT Tap captures an existing session through a mirror port. It sends nothing to either endpoint.

The monitor needs the listener endpoint and, for encrypted content analysis, the session passphrase. When a listener serves several callers, the caller's address identifies the session to monitor.

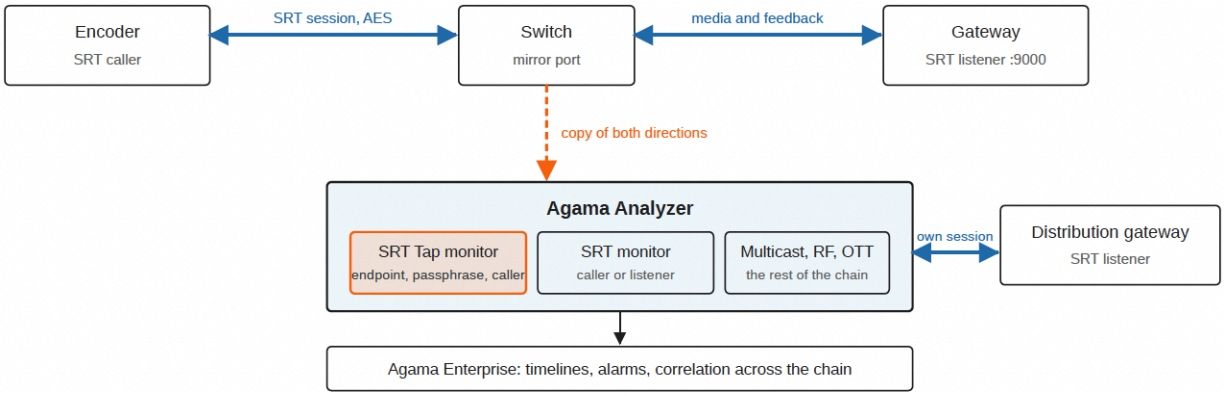

Figure 3. The Analyzer passively monitors the contribution session through a mirror port. It also connects to the distribution listener through a separate SRT monitoring session.

Reconstructing the played stream

SRT Tap reads the handshake to identify endpoint roles, negotiated latency and key exchange information. It also reads receiver feedback, including loss reports and round-trip timing information.

The tap decrypts data packets and models the receiver's playback decisions. It reconstructs the stream after packet recovery and deadline-based drops.

The Analyzer then applies its standard analysis to that stream:

- Transport-stream checks, including TR 101 290 and PCR analysis.
- Video and audio decoding.
- Black frame, freeze, silence and loudness detection.
- SCTE-35 analysis.
- Thumbnails and recordings.
- Quality of experience measurements and alarms.

Measuring the session

The Analyzer combines session counters with measurements of timing and loss location.

GROUP	MEASUREMENTS	OPERATIONAL USE
SRT	Receiver RTT, packet counts, loss reports, drops, decryption failures and reported loss	View the actual session's SRT statistics in a consistent format
SRT Tap latency	Negotiated latency, required latency, headroom, RTT at the tap and jitter	Assess the time available for recovery and review the latency setting
SRT Tap loss	Loss before the tap, retransmissions, loss after the tap, unrecovered loss and feedback-path loss	Locate the affected part of the path and identify loss that recovery could not repair

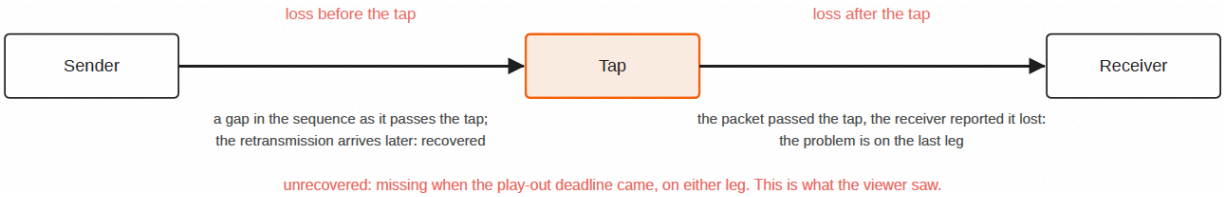

Figure 4. The tap compares captured media packets with receiver loss reports. This lets it distinguish loss before the tap from loss after it and identify packets that were not recovered in time.

Connecting transport events to content errors

The Analyzer adds SRT events to its detailed log. These include session starts and ends, handshakes, key events, packet loss, drops and decryption failures.

It also records when three conditions begin and end:

> Latency is exhausted.
>
> The feedback path is impaired.
>
> A retransmission burst occurs.

The Analyzer reports configuration issues found during the handshake once, rather than repeating them throughout the session. SRT events appear on the same timeline as transport-stream errors and video or audio faults.

For example, an operator investigating a complaint at 20:14 could see:

> A retransmission burst beginning at 20:11.
>
> Headroom reaching zero at 20:13:52.
>
> Three packet drops, a continuity error and a two-second freeze.
>
> Loss located after the tap.

This gives the operator a timeline and a narrower area to investigate.

Monitoring individual receivers

When a distribution listener serves several callers, each monitor uses a caller address to select one session. Operators can monitor each affiliate separately through the same mirror port.

For a listener with one caller, no caller filter is needed. The monitor follows reconnects automatically.

Three examples

Contribution from an event

The venue connection deteriorates during the evening. Retransmissions rise and headroom falls, while the output remains clean.

These measurements give the operator time to act. Depending on the system, they can adjust latency or switch to a backup path before the picture breaks.

Reviewing latency before replacing a leased circuit

An operator tests SRT as a replacement for a leased circuit. The initial configuration uses 200 ms of latency.

A week of measurements shows that 60 ms would have covered the recovery time observed during the trial. The operator uses this evidence to review the setting and choose a suitable margin. Handshake checks confirm the latency negotiated by the endpoints.

Investigating one affiliate’s reception problems

One of twelve affiliates reports picture artefacts. The sender’s aggregate statistics do not reveal which session has a problem.

Twelve tap monitors share the gateway’s mirror port, one for each caller. They show loss after the tap on the affected affiliate’s session. The operator includes this evidence in the support ticket.

Deployment

A mirror port can carry traffic for multiple SRT sessions. The switch must mirror both directions of the links being monitored, and the Analyzer uses a separate monitor for each session.

The monitor stores the passphrase locally and does not transmit it to the endpoints. Without the passphrase, it can still provide protocol statistics, events and configuration checks. Content analysis requires decryption.

Passive monitoring requires no changes to the SRT endpoints.

SRT monitors run alongside multicast, RF, SDI and OTT monitors. They use the same event log, alarm model and Agama Enterprise platform, so operators can compare faults across the delivery chain.

Conclusion

SRT can keep video playing while network conditions deteriorate. Output monitoring alone may miss the rising retransmissions, shrinking recovery time and feedback-path problems that precede packet drops.

The Agama Analyzer exposes these conditions. It can join a dedicated monitoring session or passively observe the session carrying the service.

With passive monitoring, the Analyzer locates loss, measures latency headroom and reconstructs the played stream. It then applies its transport-stream, video and audio analysis to that stream. Operators can investigate SRT faults in the same views and event timeline they use for the rest of the delivery chain.

About Agama

Agama is a specialist in video observability and analytics, providing an end-to-end platform for video service providers. The solution brings together insights from video creation, network delivery and viewer consumption, helping Operations, Customer Care, Product and Management teams understand service performance, investigate issues and make better use of video service data.