Standardized Status Monitoring on NMOS Systems

by Stefan Ledergerber (Simplexity GmbH), on behalf of Swiss Radio and Television (SRF)

Introduction and Management Summary

The objective of this document is to stimulate a discourse regarding an absent functionality in NMOS, informed by practical insights gathered over several years of managing a full-IP infrastructure in Zürich, Switzerland. While IP based audio/video connectivity is currently operational, SRF identifies a lack of a simplified and efficient methodology for monitoring and surveillance within their IP infrastructure. To embark on this next phase of integrating SMPTE-ST-2110 based systems into daily operations and define processes suitable for broadcast engineers for maintaining the infrastructure, SRF engaged with some key manufacturers and undertook multiple rounds of thorough discussion. A primary conclusion derived from these dialogues was the necessity to define the monitoring capabilities requested such that implementation burdens on the vendor side is kept minimal. Simultaneously, the suggested monitoring capabilities should augment the efficiency in the daily work of broadcast engineers. The minimal goal of such a monitoring system is that a broadcast engineer without deep system knowledge is enabled to detect an error in signal connections and engage the suitable 2nd or 3rd level support efficiently. This resulted in a proposed set of 4 alarms as a minimum:

1. Link Down
2. RTP Packets Lost/Late (Buffer Underrun)
3. PTP Grandmaster Change/Unlock
4. Invalid RTP Stream Format (Unable to Decode Packets)

This paper is distributed to interested vendors and standardization committees to solicit feedback and thereby expedite the implementation of a substantial yet pragmatic monitoring within Audio/Video-over-IP (AVoIP) -based products. SRF is open to alternative propositions and refinements of their proposals, provided they adhere to the two core objectives:

- Significant practical impact for daily operation
- Minimal implementation effort for manufacturers

Ideally the outcome will be included in a standard such as NMOS, and all compatible products will offer this basic set of alarms with short time to market to solve the pressing issues.

A foundational rationale of this proposal is that within a non-deterministic network (a given characteristic of standard layer-3-based networks), the only reliable points for probing the state of media streams are the end devices themselves, not the switches or any other node in between them. Only the end devices can accurately interpret the content of a packet up to the RTP payload level, and thus they are uniquely positioned to signal any problem along the whole connection path. To mitigate unnecessary traffic and computational efforts across the network and monitoring software, it's imperative that an abnormal state is proactively signaled by the edge device. This rationale underpins SRF's advocacy for active alarming as opposed to polling mechanisms.

Ultimately, it is believed that this step in media networking will support a quicker and smoother transition from traditional baseband audio/video installations to IP-based approaches. It intends to assist non-IT professionals in identifying issues within their setup, thereby fostering trust in this relatively new, still "black box" technology for many operators. Consequently, it may help accelerating its acceptance across all audio/video domains, even outside broadcast.

Current Market Landscape in Broadcast IP Transport

The open standards SMPTE ST 2110 and AES67 have become widely accepted in the broadcast industry for transporting video and audio streams over IP networks. While these standards detail transport and synchronization mechanisms and ensure a considerable degree of interoperability among vendors at a stream level, they deliberately do not specify any kind of control or setup mechanism of AVoIP devices. NMOS endeavors to bridge this gap, covering a broad spectrum of functionalities. However, only two of them seem to have achieved significant market penetration today:

- IS-04: Device Discovery and Registration, including their senders and receivers
- IS-05: Management of Connections between media nodes which are registered via IS-04

In real-world scenarios, the majority of NMOS control software presumes correct node setup by a trained user and maintaining a relatively static configuration thereafter. In addition, NMOS anticipates technical stability during and after establishing a connection. The EBU's "Technology Pyramid for Media Nodes", updated periodically, aptly represents the prevailing market scenario using open standards. The 2023 version underscores SRF's observation regarding the underdevelopment of monitoring capabilities:

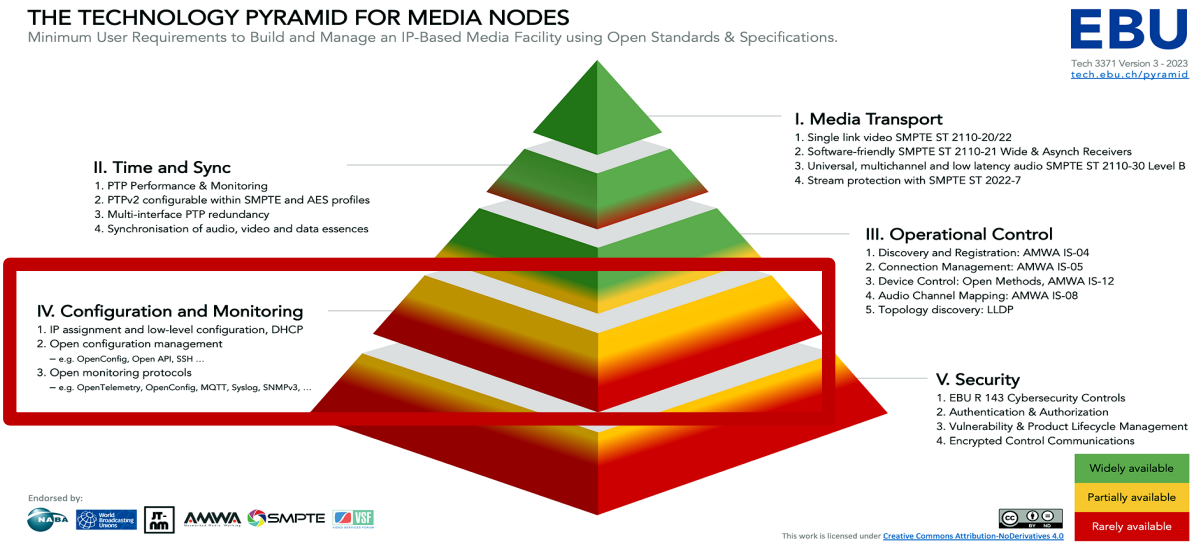

Alternatively, there are various proprietary solutions on the market for transporting audio and video over IP networks. Probably due to their single-vendor origin, these solutions often show a more rapid evolution in terms of feature completion and usability, tackling some of the illustrated challenges above more promptly than open standard approaches and often provide a better overall user experience. Prominent examples like Dante (by Audinate) or NDI (by Newtek/Vizrt).

To provide functionalities beyond basic connections, control/monitoring software, manufacturers today need to become very product-specific, using various specific control interfaces (APIs). This holds true also in media stream configuration and monitoring. Today, no unified monitoring method determines the state of a connected media node on crucial aspects like:

- Connectivity
- Synchronization
- Stream connection

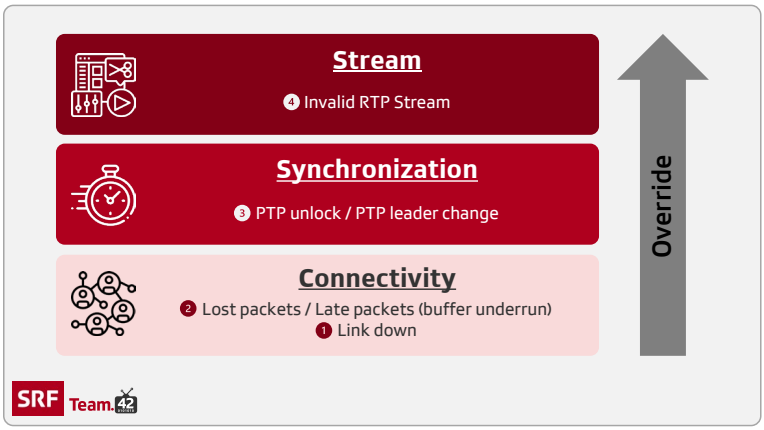

This Diagram represents a hierarchy of necessary conditions for a working AVoIP system. If a problem is identified at a lower layer, it is very likely also triggering malfunction of higher layers. E.g., if there's a problem with packet loss or a lost link, operators will also face issues all the way up to stream level. In practice, in this example, they can ignore stream and synchronization errors and focus on fixing connectivity problems, since it's most likely the root cause of all other issues.

Any problems in any one of these layers will most likely lead directly to baseband errors in media connections, unless protected by a redundant network. But even then, if a redundant installation adheres to SMPTE ST 2022-7 with dual redundant streams, there is lack of standardized mechanisms to pinpoint issues in one network (albeit invisible in baseband signals), while the other is still functional. Although some products offer some alarm features, there's an absence of standardization, and a consensus is yet to be reached on which alarms are essential for real-life operation.

Today's IP installations therefore pose challenges for broadcast engineers, who often rely heavily on pure trust into the system, without the capability to swiftly detect issues and their root causes.

Current Situation Overview at SRF

SRF has transitioned to an all-IP infrastructure that has been on-air since 2020. This project, broadly known within the industry as "Metechno", is based on the SMPTE ST 2110 standard and chose to radically move away from traditional baseband formats. In controlling media devices, the broadcast control layers use NMOS IS-04/05 wherever possible. However, to fulfill SRF's requirements including high dynamics in changing stream formats, in many cases there's a need for "hybrid" software drivers that combines basic NMOS functions with features via proprietary interfaces.

Already early in the project it became apparent that the monitoring aspects needed to be separated and delayed into a subsequent "phase 2" of the project. Presently, broadcast engineers operate the system kind of blindly, without having a satisfactory way to survey and monitor its correct operation. And even if the system would report detailed alarms, typical broadcast engineer will not be able to interpret them and derive a suitable set of action, since typical IT state messages are not targeted at end users. Only a few selected specialists possess the expertise to diagnose and address network challenges in detail. Efforts are being made to expand knowledge among colleagues, but fast scalability of such knowledge reaches its limits.

Acknowledging these areas of improvement, SRF is moving forward with the "phase 2" implementation. The aim is to simplify the system's status for non-IT professionals and to pinpoint at the relevant aspects for daily operation efficiently. A proposed set of alarms shall enable to define operational processes whereby the operator knows which department to inform, based on the alarm indicated on the system.

Steps taken so far

To optimize monitoring and system feedback, in-depth discussions were held with some key vendors to determine the best approach for SRF's requirements. Engaged in the discussions were:

- Sony (Nevion)
- Vizrt
- Pebble

From these discussions, at first long lists of potential monitoring parameters were created. However, it was soon realized that these were too extensive to be implemented in due time and system complexity may defeat the purpose of targeting the measures to typical end users. Alarms would need to be correlated and summarized to be of practical use for broadcast engineers. Examples of such detailed monitoring parameters are provided in Appendix 3. An additional challenge faced was that certain monitoring parameters couldn't be achieved with the currently installed products. Soon it became clear that for comprehensive monitoring, costly developments would be needed for some products, and maybe still not be effective. Here are some topics found difficult to monitor:

- Does any traffic arrive/leave an interface at all?
- Do the packets/traffic on a specific interface match expectation?
- Is there any packet loss happening across the entire signal path from sender to receiver?
- Is the level of packet jitter at a normal state?
- Are the buffers critically low or high (underrun/overrun)?
- Is the device synchronized correctly?
- Is the SMPTE ST 2022-7 based system fully functional or are there errors in one of the two systems which are currently being covered by the second one?

After consideration, a minimal, yet practical list of alarms was agreed upon: A minimal set of four alarms shall maximize benefits and flag issues ranging from basic hardware problems to more complex matters like decoding data streams and timing of data packets. These four alarms are associated to the layers described above, such that a control system can chose to indicate only the most significant alarm. This aspect is illustrated in the figure below as “override”. E.g. if a link is down, there is no need to display PTP unlock or missing packets.

The intention is that broadcast engineers with basic SMPTE ST 2110 knowledge will understand these alarms or are at least able to notify the correct department to fix the problem. Taking redundancy into account, it is a vital part of the proposal that these alarms are reported on an interface-by-interface basis, such that errors within redundant systems can already be detected before the second systems fails too.

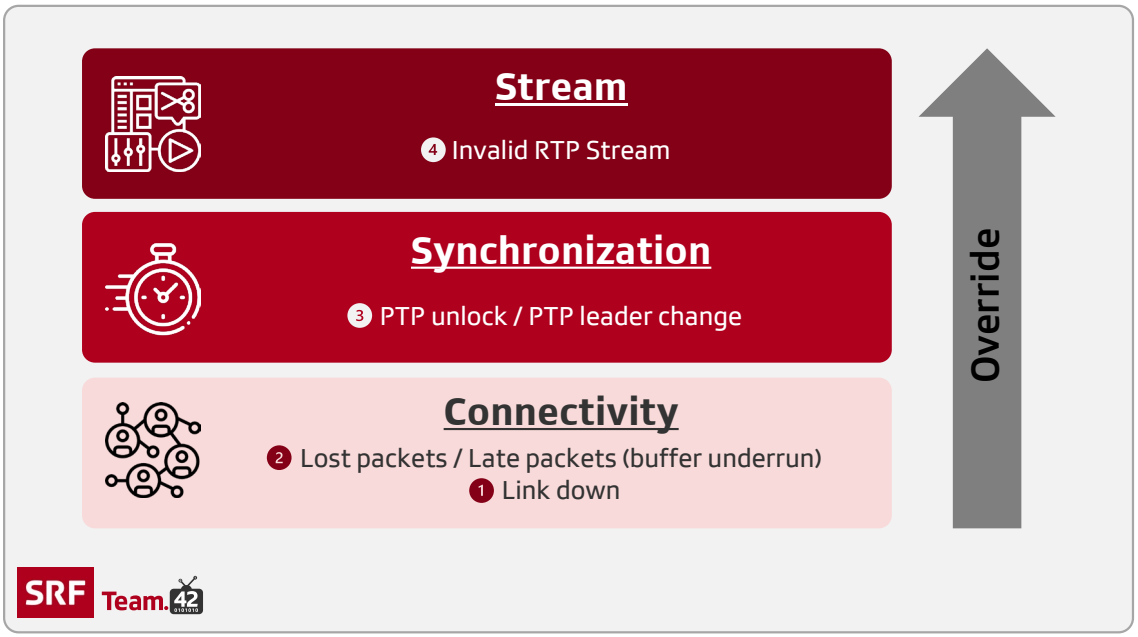

During SRF's discussions, it was noted that some essential information isn't available in some of today's products and vendors will have to do implementation work anyway. It is assumed that keeping the list of alarms short and standardize on them will help all involved parties, the equipment vendors to minimize development efforts, the users to get to know these messages and finally the control/monitoring system manufacturers to only have to implement such alarms once for all NMOS compliant devices. Since the suggested alarms are basic and clear, there's hope that this effort will work out.

As a summary, the main reasons for picking the final four alarms were:

- Covering all "layers" of a working an AVoIP system
- Spotting common problems quickly without needing IT experts and therefore helping organizations spot and fix problems fast.
- Clear definition and minimal set, therefore, reduce effort for vendors to implement meaningful alarming
- Enable organizations to implement processes which guarantee a timely solution to AVoIP problems by involving the right experts efficiently.

While deciding on the alarms was important, the group also talked about a protocol to submit them. All involved parties liked the idea of using NMOS IS-12 because of its flexibility and ability for active notifications rather than polling. But since NMOS IS-12 wasn't ready yet, NMOS IS-07 was chosen as a temporary solution. In any case we prefer an implementation following the NMOS standards.

Summary: SRF's Main Points

1. Monitoring must be done by edge devices, not by switches in between them. Only they know the expected state on a signal/stream level and can report a full picture, which takes all layers of interconnection into consideration.
2. For the system to work on a large scale, edge devices need to actively tell when something's wrong (send alarms). They should not expect a control system to poll large amounts of telemetry data.
3. Problems shall be identified on a set of "layers". This approach seems to identify true causes for problems quickly and at least involve the right department of the organization:
    a. Connectivity alarms
    b. Synchronization alarms
    c. Streams format alarms
4. Standardize a minimum set of alarms such that the user experience is the same across products, while implementation efforts on vendor side are kept minimal.
    a. Connectivity alarms
        i. Link Down
        ii. Packets lost / late (buffer underrun)
    b. Synchronization
        i. PTP unlock
        ii. PTP leader change
    c. Stream format
        i. Invalid stream format
5. Alarms should be sent for each interface separately. This way, issues in every part of a redundant SMPTE ST 2022-7 system are spotted quickly and in due time, **before** both parts show the problem.
6. Optional thought: NMOS IS-12 may serve as a base in order to foster adoption in the market.

Verification: Examples of practical problems covered by these 4 alarms
In a brainstorming session the following potential practical problems were identified. It became apparent soon that the proposed alarms would trigger in any of these cases and therefore the verification was successful. In a control system it is expected that display of the alarms follows a priority of the layers. E.g., if a link is down, the "upper" alarms can be ignored.

	Link Down	Lost/Late packets	PTP GM Change	Invalid Stream
Cable disconnect	X			
Aging SFP / dirty optical connectors / link flapping	X	X		
Incorrect packet forwarding in network (e.g. IGMP/PIM mistakes)		X		
Unexpected high traffic on network / oversubscription (e.g. IGMP timing, multiple Queriers, QoS)		X		
Link Offset (latency) in receivers set too low		X		
PTP leader degraded or lost			X	
Wrong PTP domain set by user			X	
Wrong PTP profile settings by user			X	
Incompatible stream settings by user				X
Multi-use of same multicast address				X

Proposed Next Steps
This document is intended for broad dissemination within the broadcasting industry. We advocate for a recognized standardization body to spearhead this distribution and take over the initiative. Feedback is eagerly anticipated on the following topics:

- The viability of the proposed alarms
- Feedback on the recommendations made
- Overall reflections on the document
- What does it take to implement these alarms?

Our aspiration is that this initiative will stimulate further collaborative efforts throughout the broadcasting sector. The hope is that these requirements resonate not just with SRF but are universally acknowledged across the industry. Ultimately, we envisage that a large number of vendors affiliated with NMOS will integrate these alarm protocols and NMOS alarms will become as normal as NMOS IS-04/05.

Contact

SRF

Swiss Radio and Television SRF
Sandro Furter, Senior Project Manager
Fernsehstrasse 1-4
8052 Zürich
Switzerland

sandro.furter@srf.ch
team.42@srf.ch

www.srf.ch

Author

Simplexity GmbH
Stefan Ledergerber, CEO
Schärenmoosstrasse 78
8052 Zürich
Switzerland

stefan.ledergerber@simplexity.ch

www.simplexity.ch

Appendix 1: Specification of alarms

Alarms may include optional details like severity, time of occurrence, and set thresholds which are not mandated below. Such thresholds may be adjustable through an API, but the manufacturer must set them to a relevant default value. Clients can subscribe to specific alarms. However, it can stop these subscriptions if they are deemed irrelevant for a particular scenario or user context (referred to as "alarm mute"). For instance, if a 'link down' alarm is activated, the monitoring system might unsubscribe from all related higher-level alarms concerning interrupted connections through that specific link.

User Story 1:
While using a fully redundant A/B network, I've set up a connection in SMPTE 2022-7 format. If a link in network B accidentally disconnects, my signal remains uninterrupted. I expect my monitoring system to alert me with a "link x down" warning so I can address it during my next available break. Ideally, the system would also inform me about which signals might be at risk due to this issue, namely those that pass through the affected link and are currently without backup.

User Story 2:
Following the scenario in User Story 1, if an additional link in network A also disconnects, resulting in the disruption of one or multiple signals, I expect my monitoring system to issue a critical "link y down" alarm. Ideally, it should also indicate which signals have been affected. Since the main issue stems from the base layer "connectivity," there's no need to alert me with higher level alarms 2-4, which all originate from the same root problem. I already know that I have an issue with the network and will call the IT department. To avoid confusion, I want my monitoring system to refrain from notifying me about these subsequent alarms that naturally arise due to the "link down" situation.

1. **Link Down**
    1. Alarm condition: A link was established and then gets lost
    2. Alarm clearance: Acknowledge by user or monitoring system (API) or the link is re-established again.
    3. Report interface in question
    4. (In case of in-band-management, alarm cannot be reported)

2. **RTP packets missing or late**
    1. Alarm condition: An active RTP receivers did not have media content in buffer at the time of expected playout time (link offset in AES67)
        a. Start alarm detection 5 seconds after a new connection is established
        b. Ensure alarm is not triggered upon disabling a receiver (user action: disconnect)
        c. Don't alarm about inactive receivers
    2. Indicate number of packets which did not arrive in time (late packets, buffer underrun)
    3. Indicate number of packets which were not received at all (lost packets)
    4. Report as a counter
    5. Alarm Clearance: Reset of counter by user or monitoring system (API) or receiver disable
    6. Report for both redundant interfaces separately

3. **Synchronization**
    1. Alarm condition: An interface is locked to PTP and then changes its state to either another GM or unlock
    2. The new state may be reported as "unlocked" or the new GMID, completed with previous GMID/state.
    3. Alarm clearance: Acknowledge by user or monitoring system (API) or the interface is locked to the previous GM again.
    4. Report for both redundant interfaces separately

5. In case of SMPTE ST 2022-7 redundancy, this implies that both interfaces are constantly monitored for PTP traffic, even the passive one.

4. Invalid RTP stream format

1. Alarm condition: An RTP receiver is enabled and does not recognize the stream format or is not able to decode the RTP packets
    a. Start alarm detection 5 seconds after a new connection is established
    b. Make sure alarm is not triggered upon disabling a receiver or user action “disconnect”
    c. Don’t alarm about inactive receivers
2. Alarm clearance: Acknowledge by user or monitoring system (API) or receiver is able to decode the stream

Appendix 2: List of discussed telemetry and alarms

The following list of possible alarms and telemetry states have been part of discussion (extract) but have been reduced to the proposed 4 alarms.

	Telemetry	Severity	Comment
RTP			
Stream Standard	Shall		
Stream Format	Shall		
Invalid Stream Format	Must	Alarm	
Late Packets Interface 1	Must	Warning	
Late Packets Interface 2	Must	Warning	
Missing Packets Interface 1	Must	Warning	
Missing Packets Interface 2	Must	Warning	
Packet Interval [ns] Interface 1	Must	Warning	= Jitter
Packet Interval [ns] Interface 2	Must	Warning	
Recovered Packets using second link	Must	Warning	How to reset? Counted against Missing Packets? (13)
SPF			
Link Active	Must	Alarm	tested
Bit Rate Nominal [Link Speed]	May		
Transmit rate [bytes/second]	Must	Warning	If more than 90% of total bandwidth used
Receive rate [bytes/second]	Must	Warning	If more than 90% of total bandwidth used
Transmitted packets	Shall		
Received packets	Shall		
Dropped transmit packets	Shall		
Dropped receive packets	Shall		
Transmit Output Power [uW]	Shall		Optional (Difficult to identify the limits) Not supported by VIF
Receive Input Power [uW]	Shall		Optional (Difficult to identify the limits) Not supported by VIF
Link Length SMF [km]	May		
Link Length SMF [100m]	May		
Link Length OM2F [10m]	May		
Link Length OM1F [10m]	May		
Link Length OM4/Copper [10m/mkm]	May		
Link Length OM3 [10m]	May		
Temperature [Celsius]	Shall		Optional (Difficult to identify the limits) Not supported by VIF
Supply Voltage Transceiver [mV]	Shall		Optional (Difficult to identify the limits) Not supported by VIF
SFP Label	May		
Vendor Name	May		
Transceiver Type	May		
Connector Type	May		
PTP			
Current PTP state [locked/unlocked]	Must	Alarm	e.g. locked/unlocked
Best Master Clock Id	Must	Warning	upon change of GMID
Grandmaster Clock Id	Must		
Clock Offset [ns]	Shall		
Path Delay [ns]	Shall		