



NRK Device ST 2110, ST 2059 and NMOS Conformance

601 Platform Standards

1. Contents

1. Contents.....	1
2. Version Control.....	2
3. Introduction.....	3
4. Platform Configuration Overview.....	4
5. Standards.....	6
6. Exceptions to EBU Tech 3371: Minimum User Requirements to Build and Manage an IP- Based Media Facility Using Open Standards and Specifications.....	7
7. Exceptions to JTAM TR-1001-1:2020 System Environment and Device Behaviours for SMPTE ST 2110 Media Nodes in Engineered Networks.....	9
8. Interoperability.....	10
8.1. RTP Payload Values.....	10
8.2. Synchronisation, Timing and Latency.....	10
8.3. Traffic Shaping.....	10
8.4. Receiver WAN Interoperability.....	11
8.5. Audio Devices and Endpoints.....	11
9. Device Behaviours.....	12
10. Control.....	14
10.1. Device Performance.....	14
10.2. NMOS Resource Labelling.....	14
10.2.1. Requirement.....	15
10.3. Multiformat.....	15
11. Security.....	16

2. Version Control

Rev.	Date	Detail	Author(s)
3.0	17/06/22	First draft release of document for review Alignment with architectural model and AVoIP live media production formats.	GH
3.1	27/06/22	Updates from NRK review: Stig Krokstad, Henning Bernsten, Thomas Berglund, Odd Erling Høgberg, Erik Vold, Dag Gulbrandsen. 1. Added NMOS IS-07 and 09 to optional standards support table. 2. Added mandatory RTP Offset = 0 parameter to ensure AES67 / ST 2110 compatibility. 3. Added requirement for 64 channel flow for certain categories of audio devices.	GH
3.2	17/08/2022	NMOS IS-07, 08 and 10 made mandatory in Section 5 standards tables. Updated link to TR-1001-1. Section 7. Update to IS-04 heartbeat support parameter support. Added section 11. Security, requiring demonstration of robust cyber security practices per EBU r143. Updated diagram in Section 4.	GH
3.3	02/09/2022	Section 9. Point 13 - Added requirement. «It must be possible for a Receiver to subscribe to a multicast group that originates on the same interface (i.e. loopback).» Section 9. Point 8 – Minor adjustment to expected device behaviour.	GH
3.4	10/10/2023	Introduced NMOS IS-11, 12, 13, IS-06, IS-08, IS-11, IS-12 Updated wording NMOS resource labelling wording Made some adjustments to audio requirements relating to the carriage of uncompressed and compressed audio within the same flow.	GH

3. Introduction

To support the transition to AVoIP production, our industry has created many new standards and built on many existing ones.

To build supportable, interoperable IP production platforms, it is critical that vendors implement these standards in a consistent way. Furthermore, it is necessary that vendors support the specific configuration options within those standards necessary for successful deployment within NRK's AVoIP infrastructure.

This document details the NRK's conformance requirements for AVoIP devices supporting ST 2110, ST 2059 and AMWA NMOS standards.

Broadly speaking, endpoints must adhere to the requirements detailed in EBU Tech 3371 (<https://tech.ebu.ch/docs/tech/tech3371.pdf>), AMWA NMOS JT-NM_TR-1001 (<https://www.jt-nm.org/tr-1001-1>) and AMWA NMOS BCP-003 (<https://amwa-tv.github.io/nmos/>). Additionally, endpoints must adhere to the requirements detailed in this document.

When reviewing the documentation, vendors are requested to state their product compliance by marking the feature "comply" and where more detail is required, describe the current product capability. Where a feature is not currently available within a product, vendors must state their current capability and when the required feature will become available. Where a feature is not relevant to a product, vendors should simply mark it "not applicable". Where a feature will be implemented vendors should mark it "roadmap [delivery date]". Where a vendor has no plans to implement a requested feature, it should be marked "does not comply".

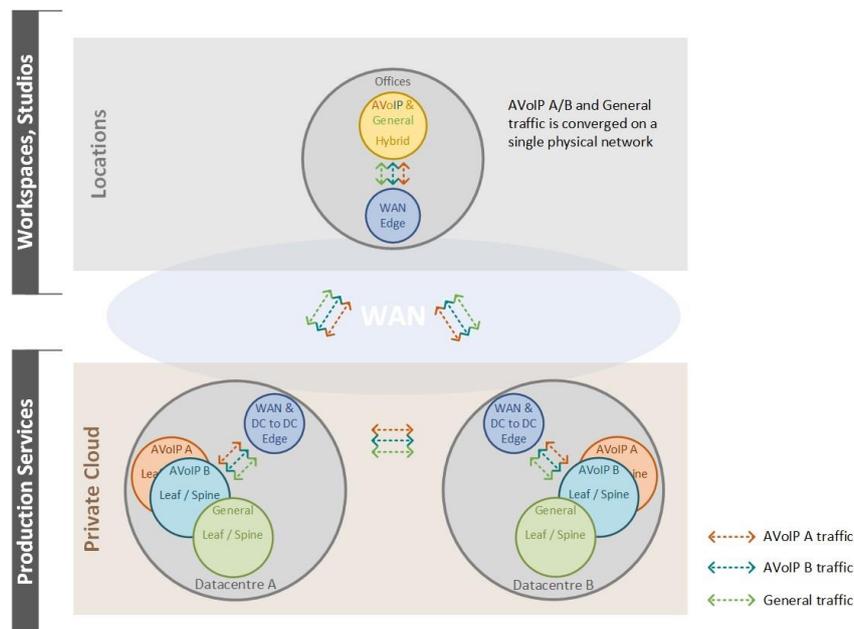
4. Platform Configuration Overview

NRK's production platform shall be built around two datacentres.

Within each datacentre is all the backend processing equipment required to support operational workspaces (e.g., edit suites, PCR, MCR, Tx suites etc.) located in different regional offices in Norway.

Regions and datacentres are all connected via a nationwide, high-bandwidth WAN. Additionally, datacentres are directly connected via resilient and diverse high-bandwidth network connections.

Within a datacentre, two independent leaf/spine networks (Fabric A and Fabric B) provide the audio/video switching fabric. A third leaf/spine network provides control and monitoring connectivity. Devices supporting ST 2110 are connected to both A and B networks. Resilient, GPS locked PTP GMs provide ST 2059 profile PTP to network fabrics. Each leaf will be configured as a Boundary Clock (BC). BC's shall select the best clock to lock to. Endpoints shall lock to the local BC.



Within a region, control and AVoIP traffic is converged on a single network. The network is virtualised and dedicated AVoIP uplinks between switches create A and B networks and carry ST 2110 traffic. Resilient, GPS locked PTP GMs provide ST 2059 profile PTP to network fabrics. Each leaf will be configured as a Boundary Clock (BC). BC's shall select the best clock to lock to. Endpoints shall lock to the local BC.

Devices deployed across the platform are required to be interoperable across the WAN and between datacentres. The instantaneous path differential (PD) timing for resilient ST 2022-7

flows between the datacentres shall be less than 10ms. Between a datacentre and region, instantaneous path differential timing for resilient ST 2022-7 flows shall be less than 50ms. Receivers must be 2022-7 Class B compliant and be able to tolerate an instantaneous PD of up to 50ms.

Network flows within the DCs and Offices are orchestrated by SDN controllers. No other SDN controllers are permitted within the platform.

All endpoints must support IS-04 for registration/discovery and IS-05 for direct device control. Each datacentre and office shall have a separate IS-04 registry.

5. Standards

Devices must support the following standards, where applicable:

#	Standard	Title
1	SMPTE ST 2059	SMPTE Profile IEEE-1588 PTP
2	SMPTE ST 2110-10	System Timing and Definitions
3	SMPTE ST 2110-20	Uncompressed Active Video
4	SMPTE ST 2110-21	Traffic Shaping Uncompressed Video
5	SMPTE ST 2110-22	Constant Bit-Rate Compressed Video
6	SMPTE ST 2110-30	PCM Digital Audio
7	SMPTE ST 2110-31	AES3 Transparent Transport
8	SMPTE ST 2110-40	SMPTE ST 291-1 Ancillary Data
9	SMPTE ST 2022-7	Seamless Protection Switching
10	AMWA NMOS IS-04 v1.3.x	Discovery and Registration
11	AMWA NMOS IS-05 v1.1.x	Device Connection Management
12	AMWA NMOS IS-07 v1.0.x	Event & Tally
13	AMWA NMOS IS-10 v1.0.x	Authorization
14	AMWA NMOS BCP-003	Security recommendations for NMOS APIs
15	IEEE 1588v2	Precision Clock Synchronization Protocol for Networked Measurement and Control Systems
16	IEEE 802.1AB	Specifically, the Link Layer Discovery Protocol (LLDP)
17	AES67	High-performance streaming audio-over-IP interoperability
18	IETF RFC 3376	IGMPv3/SSM
19	ISO/IEC 21122	JPEG XS low-latency lightweight image coding system
20	AES-r16-2016	PTP parameters for AES67 and SMPTE ST 2059-2 interoperability
21	VSF TR-08	Transport of JPEG XS Video in ST 2110-22

Table 1. Mandatory standards support

Devices should support the following standards:

#	Standard	Title
1	AMWA NMOS IS-08 v1.0.x	Audio Channel Mapping
2	AMWA NMOS IS-09 v1.0.x	System Parameters
3	AMWA NMOS IS-11	Stream Compatibility Management
4	AMWA NMOS IS-12	Device Control
5	AMWA NMOS IS-13	Annotation
6	SMPTE ST 2110-41	Fast Metadata [when it is released]

Table 2. Desirable standards support

6. Exceptions to EBU Tech 3371: Minimum User Requirements to Build and Manage an IP-Based Media Facility Using Open Standards and Specifications

Devices must adhere with the requirements detailed in the latest version of EBU Tech 3371, with the following exceptions:

- I.1.1 Single link video: SMPTE ST 2110-20

*“For simplicity of operation, troubleshooting and optimally dense use of the network, video Media **shall** use single link streams (i.e. Media Nodes should not use SMPTE ST 2110-23). For instance, a single 59.94 Hz UHD stream requires a 25 Gigabit Ethernet (GbE) port and a high-density device (such as a multiviewer or a vision mixer) **shall** take benefit of bi-directional 100 GbE.”*

Please note alterations to text above.

- I.3.1 Universal, multichannel, low latency audio: SMPTE ST 2110-30 Level B

“In addition to SMPTE ST 2110-30 Level A, which is compatible with all AES67 and 2110-30 devices, audio Senders and Receivers shall support Level-C to support low latency applications and future multi-channel audio applications;”

Please note alterations to text above.

- I.4.2 Stream protection with SMPTE ST 2022-7:2018

“Receivers shall support SMPTE ST 2022-7:2018 Class D with Ultra Low-Skew 150 μ s that is adapted to engineered LAN with adding minimal latency”

Requirement does not apply.

- II.2.1 & II.2.2 PTPv2 configurable within SMPTE and AES profiles

“PTP parameters of Media Nodes shall be configurable within the union of the ranges covered by the both SMPTE ST 2059-2 profile, the AES67 media profile and the IEEE-1588 default profile so that we can use the Media Nodes for the full range of possible operational scenarios;”

“However, it is recommended to the users to select an operational point within the range of AES-r16-2016 so that it is compatible with equipment previously delivered on the market that might be limited to one of the two profiles;”

Please ignore. PTP GMs shall be configured with SMPTE ST 2059-2:2021 default parameters. Importantly, in this release of the standard, portDS.logAnnounceInterval

has been change from -2 to 0, reducing the frequency of Announce message which enables a common PTP domain for AES67 and ST 2110 devices.

7. Exceptions to JTNM TR-1001-1:2020 System Environment and Device Behaviours for SMPTE ST 2110 Media Nodes in Engineered Networks

Devices must adhere with the requirements detailed in the latest version of JTNM TR-1001-1:2020, with the following exceptions:

- 10.2.6 NMOS System Parameters (IS-09)

*“Media Nodes ~~shall~~ **should** use the registry heartbeat_interval value specified in the System API defined in AMWA IS-09 when maintaining their registration.”*

Please note alterations to text above.

- 10.3 Multicast Media Streams

*“Media Nodes ~~should~~ **shall** provide a user mechanism for selecting the DSCP markings of the generated streams.”*

Please note alterations to text above.

- 11.2 Media Node Startup and Multicast Addresses

“The System API defined in IS-09 includes a System ID, which shall be assigned uniquely (e.g. randomly) in each facility. Media Nodes shall store the System ID and their DHCP-assigned address as part of their current operating settings, and shall check at re-start, before generating any multicast outputs, that the current Network Environment’s System ID and DHCP-assigned address match the previously stored values. If the current system ID and address are the same as before, then recalling the previous operating settings might be appropriate and safe, subject to the clause below.”

Products should support IS-09.

8. Interoperability

8.1. RTP Payload Values

1. Devices must support SDP configuration.
2. Receiver endpoints must autodetect and select RTP payload values from the Sender SDP file.
3. Devices must support SDP Payload Type values 96 to 127. See Section 6, Table 4 of <https://datatracker.ietf.org/doc/html/rfc3551> and Section 6.2 of ST 2110-10.

8.2. Synchronisation, Timing and Latency

1. As defined within ST 2059, all devices shall support IEEE1588-2008. IEEE1588v1 is not supported.
2. Default values within ST 2059-2:2021 shall be used.
3. Devices shall support a hybrid PTP communications model. Sync and Follow_Up (multicast), Delay_Request and Delay_Response (Unicast).
4. Devices must be capable of handling asynchronous baseband video and audio inputs. If alignment cannot be achieved, Audio is to be passed through out of sync and an alarm raised.
5. Devices must have an optional “minimum latency” mode, whereby baseband outputs are presented without further frame synchronisation.
6. SMPTE ST 2110-10 mandates that the media clock and the network timebase shall share the SMPTE ST 2059-1 epoch, with an offset of zero between the Media Clock and the RTP Clock, as specified in SMPTE ST 2110-10. All audio devices (specifically AES67) shall adhere to this.

8.3. Traffic Shaping

1. Senders shall adhere to ST 2110-21 ‘Narrow’ or ‘Narrow Linear’ definitions.
2. If a Sender complies with the ‘Wide’ definition, vendors must state how ‘Wide’ the Sender is.
The SMPTE ST 2110-21 ‘Wide’ definition permits a range of performance levels. For example, whilst one Sender implementation may come very close to meeting ‘Narrow’ performance, another may meet the widest definition of ‘Wide’. Senders marginally outside of the Narrow definition may (subject to testing be acceptable).
3. Receivers must be ST 2110-21 ‘Type A’ compliant i.e. capable of receiving signals from a type N, NL, or W sender, regardless of the value of the sender’s `ts-refclk` `clksrc` or the sender’s `TROFF` parameter value. Vendors must state any device performance implications resulting from supporting ‘Wide’ flows.

8.4. Receiver WAN Interoperability

1. Some endpoints will be required to receive 2022-7 protected flows which have originated from a Sender located at another site on NRK's WAN. The instantaneous path differential timing between diverse network paths across the WAN may be greater than 10ms. Vendors must state their ST 2022-7:2018 receiver class compliance.

8.5. Audio Devices and Endpoints

1. The NRK audio architecture is based around the exchange of 1-channel, 2-channel, 8-channel and 16-channel, 125us ST 2110-30 flows (Conformance Level: C). Audio Senders and Receivers responsible for processing uncompressed PCM audio must be capable of processing flows in these formats.
2. Audio devices required to send or receive flows via trunk links shall support up to 64 channels of audio per flow.
3. The NRK audio architecture also supports the exchange of 2-channel 125us ST 2110-31 flows (Conformance Level : C). Audio Senders and Receivers responsible for processing AES3 signal must be capable of processing flows in this format.
4. An audio Sender which is capable of handling more than 2 channels of audio should be capable of generating multiple 2 channels flows up to the maximum number of channels the device can handle.
5. An audio Receiver which is capable of handling more than 2 channels of audio per flow should be capable of receiving audio as 2110-30 (PCM) or 31 (compressed) within the same flow.
6. Senders should be capable of combining any mixture of uncompressed and compressed audio feeds in the same ST 2110-30 source flow.
7. For a contribution decoder the propagation delay for each Transport Stream through to ST-2110-30 sender must be consistent, irrespective of whether the audio is PCM or compressed.
8. Receivers capable of receiving multiple 2110-30 multicast groups must be able to receive groups with different channel counts.
9. Receivers capable of receiving multiple 2110-30 multicast groups should be able to shuffle audio as required. Associated with this, support for AMWA IS-08 should be planned as a roadmap feature.
10. All AES67 / ST2110 audio devices shall support 2022-7 on two physical network interfaces.
11. The device must guarantee phase coherence to within 0.2 samples between channels in a flow.

9. Device Behaviours

1. Devices must be capable of locking to PTP on both Fabric A and Fabric B network interfaces.
2. Devices must compare PTP sources using the Best Master Clock Algorithm (BMCA) defined in IEEE 1588-2008 to compare clocks on Fabric A and Fabric B and lock to the best clock.
3. In normal operation, the clock that an endpoint locks to on each of its network interfaces will be equal. In the event of a tie, the device shall use the BMCA algorithm defined in IEEE 1588-2008 to determine which interface to take time from.
4. Should a device lose lock to the BC it is currently syncing to, it must transparently failover to the PTP source on its other network interface and be capable of failing back when service is restored. This process must have zero impact on the performance of the endpoint from a user perspective.
5. A device which has lost connectivity to one fabric for a prolonged period (e.g. due to a leaf failure) will receive many routing requests during the outage, which it won't be able to fulfill on the failed link. On recovery of a failed network interface, a device shall re-acquire the current set of multicast groups to which it is subscribed. Historic subscription requests which were requested during the outage, and which have been superseded by newer requests shall be discarded. Historic route requests which have not been superseded shall be automatically acquired.
6. The possibility exists that a fault may arise which causes an endpoint to lose one flow from a logical group of flows (e.g. the audio, video or ancillary) on one or both network connections.
 - a. If the route is still current, an endpoint must re-acquire the multicasts transparently when the flow(s) are reestablished.
 - b. The loss and recovery process must not affect the stability of the endpoint. E.g. loss of the video flow shall mean that the endpoint shall display optionally black, freeze frame or another still image with the audio and any ancillary data continuing.
 - c. Loss of an audio stream shall mute relevant channels, until such time as the flow is restored. Whilst audio must be restored as quickly as possible, this must not result in loud clicks or pops as the audio is reacquired.
7. The possibility exists that an endpoint may become flooded with multicast traffic. Once the flood is resolved, the endpoint must remain stable and recover all current multicast groups it should be subscribed to without user intervention.
8. Devices are not permitted to create a 1:1 relationship between a multicast receiver and a device input. NRK may wish to present the audio within a multicast to multiple device inputs simultaneously. It is not possible for a network interface to subscribe to the same multicast more than once. An endpoint which is capable of receiving multiple flows and presenting them to an application must also be capable of internally routing the same IP input stream to all application inputs. In

such a situation, all receivers must report the 'actual' sender, rather than an any internally replicated source which may be used to work around the limitations of multicast.

9. Endpoints must reliably exchange SMPTE ST 2110 compliant flows with endpoints from other manufacturers.
10. Devices must not place any restrictions on the multicast group address assignment. e.g., restricting to 239.x.x.x.
11. Devices must allow the TTL parameter to be user configurable.
12. Devices must support SSM.
13. It must be possible for a Receiver to subscribe to a multicast group that originates on the same interface (i.e., loopback).

10. Control

1. Devices must be directly controllable via an NMOS API, rather than via separate API gateway server.
2. Devices shall support NMOS IS-04 v1.3.x and IS-05 v1.1.x.
3. Devices shall support NMOS IS-07 v1.0.x.
4. Devices shall implement the BCP-002 `GroupHint` tag.
5. Devices shall support BCP-003 to provide secure communications.
6. Where in-band control is supported, devices shall allow NMOS control on both A and B NICs simultaneously. Binding NMOS control to one NIC or the other is not permitted.
7. NMOS UIDs must not be derived from replaceable device components. For example, replacing a NIC should not result in the generation of new UIDs.

10.1. Device Performance

A logical source may consist of multiple Senders associated with multiple Devices / Nodes (typically a video, multiple audio and multiple data flows) and a logical destination may optionally select from flows from multiple Senders associated with multiple Devices / Nodes. In a ST 2110 based platform, a lot of messaging between systems must take place between systems to establish connectivity between a logical source and logical destination. To avoid unacceptably slow connection request response times, it is therefore important that control systems and devices process commands as quickly as possible and communicate with other systems efficiently.

Though other protocols are supported (RTP and MQTT), in practice the majority of NMOS implementations use the WebSockets protocol for message exchange. WebSocket connections are slow. Establishing a connection can take ~200ms and requests ~80ms to complete. It is therefore easy to see how (without careful optimisation) subscribing a device to multiple audio, video and data flows could take seconds to complete and a salvo to a multiviewer, vision mixer or audio mixer might take significantly longer.

1. Devices shall provide support for IS-05 “bulk” connection requests.
2. Devices shall process individual connection requests within a bulk connection request in parallel.
3. The Supplier shall state how long it takes a Receiver to process a bulk request containing ten connections.

10.2. NMOS Resource Labelling

The NMOS IS-04 specification defines description and label properties of nodes, devices, senders and receivers but does not provide any guidance on how they should be populated.

Consistent labelling is critical to device commissioning and ongoing platform support. For example, having 100 decoders all advertising with device labels that contain the model number

is unhelpful. Similarly, sender/receiver labels which are manufacturer specific e.g. “Sender Video 0 of Receiver device SNXXXX” will further add to the confusion.

Not all vendors implement logical grouping tags (and those that do use a variety of different syntaxes).

MPP plans to adopt a simple naming scheme for devices and senders/receivers which results in consistency across control and configuration platforms. Whilst this exact structure has yet to be formalised, the scheme might be as follows

Each device has a unique id (UID) - in the form 12345/THING/001 - the first segment denotes a physical or operational area, the second the device category (e.g., IRD/CODEC/DECODER) and the last segment is a counter. This device UID forms the basis of a sender or receiver id:

12345/THING/001_s01_v01	12345/THING/001_s01_a05
12345/THING/001_s01_a01	12345/THING/001_s01_a06
12345/THING/001_s01_a02	12345/THING/001_s01_a07
12345/THING/001_s01_a03	12345/THING/001_s01_a08
12345/THING/001_s01_a04	12345/THING/001_s01_m01

These 10 senders would represent one video, eight audios and one metadata flow - i.e., a logical grouping or spigot. For receivers, we substitute _s for _r

10.2.1. Requirement

Vendors MUST provide a mechanism for project engineers to enter the following [minimum] metadata set for IS-04 as freely definable text.

- Node label
- Device label
- Sender label
- Receiver label

This MUST be achieved through one of the following methods (in order of preference)

- Manufacturer’s API
- Plain text (e.g., xml, json) File loaded into your system
- Direct entry through web page or Manufacturer’s configuration tool

Ideally, the solution SHOULD NOT involve manual labelling of each resource

It may be permissible to employ sender/receiver patterns, such that {Node}, {Device}, {Spigot}, {Flow Type} variables can be used to build a hierarchical label definition.

Vendors are encouraged to actively participate in AMWA NMOS, as this topic is likely to form part of a forthcoming activity.

10.3. Multiformat

One of the key benefits to broadcasters of the transition to IP is that the data transmission technology (the professional media network) is format agnostic, allowing rapid adoption of new

media formats and multiple media formats to coexist on the same network fabric. Such flexibility requires care be taken when defining platform control system design and philosophy.

Increasingly, devices are becoming reconfigurable to be able to create and consume a broad range of compressed and uncompressed media formats, each with very different bandwidth requirements. Such flexibility presents challenges to platform management.

To ensure the correct operation of a large professional media network, it is critical that the opportunity for unsolicited flows to get onto the network is minimised. Furthermore, it is important to be able to detect and set correct Sender and Receiver format configuration.

1. The device shall provide an API call to report and set Sender and Receiver format configuration.
2. The Supplier shall explain how the device presents Senders and Receivers via NMOS when Sender / Receiver format/bandwidth is changed and how this relates to the internal processing capabilities of the device. For example, when switching between 1080i25/1080p50 and 2160p50 modes, several different behaviours have been observed, including:
 - a. The device presents different sets of NMOS UIDs for UHD and 3G modes.
 - b. Existing NMOS UIDs are retained but three out of four UIDs are invalidated.

11. Security

Suppliers must demonstrate that processes are in place to:

- a. Minimise the risk of product / service vulnerabilities and;
- b. Minimise the risk that of professional or support services exposing NRK to threats

Suppliers are expected to review EBU r143 and complete columns A and B of the R 143 Security Controls Assertion.

Supplier must deliver a software bill of materials.