OPERATING EUROVISION AND EURORADIO

TECH 3361-3

SERVICE LEVEL AGREEMENT FOR MEDIA TRANSPORT SERVICES

TECHNICAL PARAMETERS

Geneva
September 2014

Page intentionally left blank. This document is paginated for two sided printing

Contents

1. Introduction 5
 1.1 Purpose 5
 1.2 How to use 5
2. Common Functional Parameters 6
 2.1 Service Technology 6
 2.2 Capacity 6
 2.3 Latency (or time parameters) 6
 2.4 Technology-specific parameters 6
 2.5 Types of protection 7
3. Service Level Objectives 8
 3.1 Fault/performance quantification model 8
 3.2 The G.826 parameters 9
 3.3 Determination of the Availability/Performance status of a connection 10
 3.4 Presentation of Availability/Performance parameters 11
4. Performance Parameters 11
 4.1 Media Services - Video 11
  4.1.1 Interfaces 11
  4.1.2 ES, SES definitions 12
 4.2 Media Services - Audio 12
  4.2.1 Interfaces 12
  4.2.2 ES, SES definitions 13
 4.3 Transport Performance Parameters and impact on Media Quality 13
 4.4 Transport Services - IP 13
  4.4.1 Interfaces 13
  4.4.2 Traffic Performance Parameters 13
  4.4.3 ES, SES definitions 15
 4.5 Transport Services - Ethernet (Data Link Layer) 15
  4.5.1 Interfaces 15
  4.5.2 Performance Parameters 16
  4.5.3 ES, SES definitions 16
 4.6 Transport Services - SDH / SONET (Data Link Layer) 17
  4.6.1 Interfaces 17
  4.6.2 Performance Parameters 17
  4.6.3 ES, SES definitions 18
 4.7 Transport Services - Structured WDM and OTN (Physical Layer) 18
  4.7.1 Interfaces 18
  4.7.2 Performance Parameters 19
  4.7.3 ES, SES definitions 19
 4.8 Transport Services - Dark Fibre & bit transparent WDM (Physical Layer) 20
  4.8.1 Interfaces 20
  4.8.2 Performance Parameters 20
  4.8.3 Unavailability definition 21
5. Conclusion 21
Annex A: Path protection architectures 23

Page intentionally left blank. This document is paginated for two sided printing

Service level Agreement for media Transport Services

Technical Parameters

EBU Committee	First Issued	Revised	Re-issued
FNS-SLA	2014		

Keywords: SLA, Service Level Agreement, BMO, Broadcast Media Organisation, Media Transport.

1. Introduction

With the rapid growth of IT infrastructure, reduced provisioning cost and turnaround time, more and more Broadcast Media Organisations (BMO) are considering the use of data network services as a media transport method. However, when it comes to preparing Service Level Agreements (SLA) between BMOs and Network Service Providers (NSP), many organisations find themselves lacking knowledge in this field, especially on how to specify and negotiate the specific requirements of broadcast media organisations.

This document, prepared by the EBU Service Level Agreement (SLA) Project Group, is intended to complement EBU Tech 3361-1 Service Level Agreement for Media Transport Services - High Level Guidelines. It provides a low-level technical view of recommendations and parameters that require careful consideration when an SLA for media transport services is negotiated.

1.1 Purpose

In order to specify the technical part of an SLA, we need to define parameters for Availability and Traffic Performance objectives for the type of network service. This document outlines the key technical parameters for the main types of service that BMOs can purchase from NSPs that are associated with the Availability and Traffic Performance objectives.

The target audience for this document is technical professionals from BMOs, who contribute to the technical aspects of SLA negotiations. It is also intended to help NSPs to understand the specific requirements of broadcast media transport over their network services.

Another goal of this document is to ease the communications between the teams involved in negotiating an SLA, both between organisations but also within the same organisation, by providing a common vocabulary, as wording and concepts may vary between people with different backgrounds, such as broadcast engineers, IT specialists and telecoms operators.

1.2 How to use

This document is provided for information purposes and is intended to be used as a guideline companion to be read alongside commercial SLA documents. It may be used as a checklist to ensure important SLA aspects that are relevant to specific cases are covered. It refers to international standards that are in use in the industry, when available. It also suggests alternative metrics, for example when there is no standardized method.

The document is structured as follows: Section 2 presents functional parameters commonly used in the delivery of media services. Section 3 defines a general structure for SLA calculations. Section 4 then defines, for each type of media transport service, commonly used interfaces, the relevant performance parameters, and definitions of errors and failures for inclusion in SLA calculations.

2. Common Functional Parameters

This section describes the functional parameters common to the different types of services.

2.1 Service Technology

The technology used to provide a service may in some cases be subject to specific requirements. For example, an Ethernet service may be ordered specifically as 'Ethernet over fibre', 'Ethernet over MPLS', 'Ethernet over SDH' or 'Native Ethernet'. Likewise SDI over Ethernet/SMPTE 2022 may be a requirement.

2.2 Capacity

The capacity of a given transmission link is often also referred to as 'bit-rate', 'speed' or 'bandwidth'. This is a static parameter, which is determined in the service contract. The capacity of a transmission link should comply with the needs of the signal to be carried over the circuit, so for example an SDI signal to be carried without compression would require at least 270 Mbit/s of capacity.

When determining the required capacity of a circuit, it is necessary to take the specification of the circuit into account. For example, an IP service may be specified with an interface bit rate of 100 Mbit/s, but the actual throughput, measured as the number of bits that can be transferred per second, may be only 80 Mbit/s due to various types of internal overhead. Some NSP-provided services specify that only a certain percentage of the interface speed can be guaranteed for delivery at the other end; so for example if the speed exceeds 50 Mbit/s on a 100 Mbit/s interface, it is no longer guaranteed that the data will come through. Under such circumstances it would be necessary to order a higher nominal capacity than implied by the actual requirements of the signal. Another related factor needs be take into consideration is bearer speed. For example, the BMO maybe presented as a choice a 50 Mbit/s over a 100 Mbit/s bearer or 50 Mbit/s over a 1 Gbit/s bearer. Obviously the bigger the bearer, the more future proof it can be.

It is therefore important that the BMO specifies in the service contract precisely what net capacity or throughput and bearer are required. To make sure that the specified speed is actually delivered, the acceptance test may include a test of the net capacity of the circuit.

2.3 Latency (or time parameters)

Latency plays an important role in real-time broadcasting. For example, a good quality two-way interview normally requires latency be no more than about 50 ms. For satellite operations, it is desirable that end to end latency can be achieved within 1 s.

For broadcast signals, end-to-end latency includes physical latency as well as the latency introduced by the encoding and decoding process. Due to the jitter behaviour of bursty networks, the decoder is normally required to have a de-jitter buffer to 'smooth out' any potential artefact. This in turn adds latency.

2.4 Technology-specific parameters

Each technology has its own list of static parameters that need to be present or protocols that need to be supported by the Service. Some key examples are as follows:

AES-67 media Service: Full bit-transparency required (no bit-loss, no lossy transmission encoding etc.); Jitter on decoding / receiving side needs to comply with AES-3 standard requirements; Equal delay among separate AES-3 channels to preserve correlation where required; Preservation of the XYZ preamble sequence; Support for and preservation of 32 kHz, 44.1 kHz and 48 kHz sampling, and of arbitrary sampling frequency, etc.

IP services: TCP (RFC793), UDP (RFC768), RTP (RFC3550), RTCP (RFC3550), SIP (RFC3261), SDP (RFC4566), direct IP routing (no NATing), Support of multicast operation (IGMPv2/RFC2236, IGMPv3/RFC3376), Bi-directional traffic support (RTP / RTCP), electable port number, Symmetrical path delays (required for PTP traffic), support of PTPv2 (BC or TC, IEEE1588), Transparent support of DiffServ QoS. (refer to SMPTE Study Group on Media Production System Network Architecture if relevant), etc.

Ethernet service: Maximum Transmission Unit size (MTU), 'jumbo frames', Maximum number of MAC-addresses, Transparency to specific protocols (e.g. 802.1Q VLAN for VLAN -trunking), Multicast capability (IGMP Snooping), etc.

Dark Fibre service: Fibre Type, Transmission Distance, Polarisation Mode, Dispersion (ps/√km).

2.5 Types of protection

Service protection is an important characteristic of a media transport service. It is very important to make sure that the Service Contract specifies what type or types of protection are to be used. Often, the Service Level Agreement will indicate different availability targets according to what type of protection is deployed.

There are three primary mechanisms used to protect services in networks carrying broadcast signals:

1. Forward Error Correction (FEC)

FEC is used to protect against short-term transient disturbance to the signal. A proportion of the bandwidth of the circuit is used to carry error-correction data, generated by the encoder at the start of the circuit. This error correction data is then used by the receiver at the end of the circuit to verify the integrity of the signal and, if errors are found, to reconstruct the signal. An example schema is 'Reed-Solomon' FEC, used on ASI links, where for every 188 byte of video, 16 FEC byte are appended. It is thus known as a '204/188' schema. FEC has the benefit of protecting any circuit, whether single-path or dual, and operates synchronously with the signal clock. It can be used on any form of digital transmission system, whether packet-based or not. However, it does increase the end to end latency of the circuit, and increases the bandwidth consumed.

2. Retransmission

Retransmission is used on packet-based systems to protect against short-term disturbances to the packet flow. Mechanisms in protocols such as TCP enable the receiver at the end of a circuit to recognise that not all transmitted packets have been received, and to request retransmission of those missing. Clearly this process takes a variable amount of time, and therefore an end to end system designed to carry live-to-air broadcast signals needs to incorporate a 'de-jitter buffer' in the receive device to ensure the signal flow to the downstream system runs smoothly. The size of this buffer needs to be chosen carefully. A big buffer means a lot of retransmission can occur, but it also means the end-to-end latency of the system will be high.

3. Path protection switching

Path protection switching is used on systems which are designed with more than one path. Various techniques are used, depending on the network layer where the switching is happening, but the most common in media transport are:

- **SDH ring-resilience**, which is an integral part of SDH line-system design so does not inherently add latency. However it takes 50 ms to switch paths and thus might disturb video if FEC is not used in the IP or ASI layer carrying the video signal.
- **'Slipless' or 'hitless' switching**, which typically operates in the ASI domain and uses buffers

in the receivers to time-align two signals so switching is invisible. This is a robust solution, but does add latency to the end-to-end system.

- **'Packet merging'** (e.g., SMPTE ST 2022-7), which operates in the IP layer and merges two IP streams together to overcome upstream network disturbances. This can be used as an alternative to FEC and although double the network bandwidth is required (typically via different network routes), it eliminates the need for an FEC buffer thus reducing the end-to-end latency.

Annex A describes in more detail the relationship between path design, switching and availability.

3. Service Level Objectives

The technical part of the SLA has two major objectives, quantifying limits for:

- Availability (fault aspects of the service connection)
- Traffic performance (QoS aspects of the service connection)

Since services may be very different in their characteristics, a generic model is needed to provide a common way of quantifying fault and performance aspects of the service connection for use with SLAs.

3.1 Fault/performance quantification model

ITU-T Recommendation G.826 provides a model that defines and quantifies both the fault and performance aspects of the service connection, usually referred to as the "G.826" model.

In terms of faults, the recommendation defines the concept of "Un-available time" (UAT), used to determine service Availability. Traffic performance is defined by means of Errored Seconds Ratio (ESR), Severely Errored Seconds Ratio (SESR) and Background Block Error Ratio (BBER).

The model lends itself naturally to describe "loss" performance events, such as bit errors on an SDH connection or packet losses on an IP connection. These performance events are very easy to define in terms of CRC errors, lost packets etc. Please note it is also possible to measure performance parameters/events associated with the timing integrity of the signal, but at present there are no standardised definitions for these.

There are two types of performance event to consider; "Anomalies" and "Defects". Anomalies are pure performance events, such as randomly dropped packets. If the density of Anomalies is high, they can be integrated into a Defect. Defects can also be defined by other types of events, such as when, after cutting a fibre cable, the "Loss of Signal" Defect occurs. Receiving an Alarm Indication Signal (AIS) also constitutes a Defect detection. Defects typically cause some kind of consequent action in the system. The action could be a protection switch, for example, or it could be the transmission of a Remote Defect Indication (RDI) signal upstream. Anomalies do not cause consequent actions other than that they are counted and stored as performance data.

For packetized transport, the most important Anomaly is the discarded or missing packet. Some systems only consider this Anomaly (for example the continuity check messages (CCM) mechanism in the ITU-T Y.1731 OAM model). Then the integration of this Anomaly (for example losing 3 CCM frames in a row) constitutes a Loss of Continuity (LOC) Defect that indicates a Severely Errored Second. Other transport types (most notably SDH), have richer sets of defined Anomalies and Defects.

3.2 The G.826 parameters

The following parameters are lifted form the basis of the G.826 fault/performance quantification model:

Block: A unit of transported data. A block is a set of consecutive data bits associated with the path. For SDH/PDH the size of the block varies with the transport bit-rate; for common SDH rates a block corresponds to an SDH frame. For packet-based technologies such as IP and Ethernet, a block is simply defined as a packet or frame.

For native or baseband video and audio services, there are no established definitions for blocks in common usage, (except for MPEG-2 TS when used in DVB applications, see "ASI" section below) but it is nevertheless possible to define some. For example for an SDI service, a block could be defined as a line of the video frame. Suggested block definitions will be presented for each service technology below.

EB: Errored Block. A block that has one or more errors. Errors are typically detected by a faulty checksum that covers the block.

Note 1: When Frames or Packets are used to define the block, the absence of a block, that is, a discarded or missing frame or packet counts as an Errored Block.

Note 2: As noted above, definitions of other performance degradations that could render a block as errored, such as the block being out of a timing performance profile, have yet to be standardised.

ES: Errored Second. A one-second period with one or more Error Block (EB) or at least one Defect.

ESR: The ratio of ES to total seconds in available time during a fixed measurement interval.

SES: Severly Errored Second. A second where the ratio of Errored Blocks (EB) to delivered blocks is larger than a defined threshold, or has at least one Defect. SES is a subset of ES.

SESR: The ratio of SES to total seconds in available time during a fixed measurement interval.

BBE: Background Block Error. An errored block not occurring as part of an SES.

BBER: The ratio of Background Block Errors (BBE) to total blocks in available time during a fixed measurement interval. The count of total blocks excludes all blocks during SESs.

UAT: A period of unavailable time (UAT) begins at the onset of ten consecutive Severely Errored Second (SES) events. These ten seconds are considered to be part of unavailable time. A new period of available time begins at the onset of ten consecutive non-SES events (a non-SES event is a second that is an errored second, but not an SES, or is error free). These ten seconds are considered to be part of available time. Figure 1 below illustrates this definition.

UAS: Unavailable second. A second during UAT is considered to be an UAS.

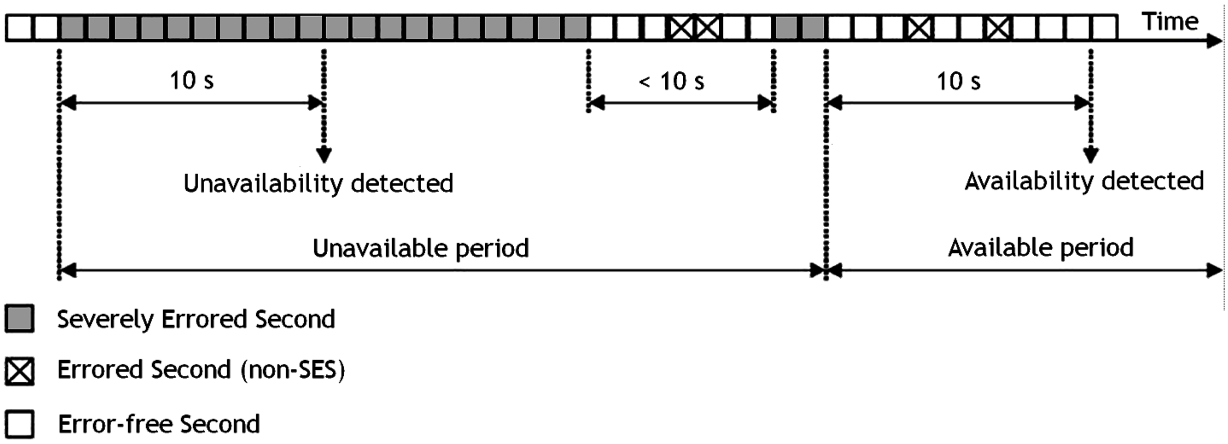

Figure 1: G.827 – Example of unavailability determination

PIU: Percent service unavailability, the fraction of the total service time that consists of UAT.

PIA: Percent service availability, 100 - PIU, the fraction of total service time that the service is available.

3.3 Determination of the Availability/Performance status of a connection

The power of the Fault Management/Performance Management (FM/PM) model is that it unambiguously provides a set of parameters describing both the availability status as well as the performance status while in service. The above parameters are easily calculated when the definition of an errored block is given. Hence the technical difficulty is confined to find these low level definitions.

Availability is associated with the UAS parameter. UAS provides UAT and UAT can be used directly in the PIU/PIA calculations, whose values correspond to the Availability criterion of the SLA.

Performance, while in service, is assessed by inspecting the ESR/SESR/BBER parameters. BBE provides a measure of how imperfect, with respect to the connection quality, the connection is. SES (while not being part of UAT) provides a measure of the level of intermittent service interruptions.

Hence the status of a connection can be divided into the following categories:

Service Availability State	Related G.826 parameters
Out of Service	**UAT:** Unavailable time (= out of service time) **PIU:** Percent service unavailability **PIA:** Percent service availability (= 1 - PIU)
In Service	**BBER:** A measure of the connection quality while in service **SESR:** A measure of intermittent service interruptions while in service

Note: There is an underlying assumption within this model that the traffic is continuous and of relatively constant bit rate, as discussed in Y.1561. This applies well to most video and audio services. For the general case, the packet loss ratio (PLR) measured as packet losses per second in relation to an established average delivered packet rate may serve to determine Severely Errored Seconds.

3.4 Presentation of Availability/Performance parameters

Performance data such as SES, ES and BBE is typically collected in 15 minute or 24 hour “bins” that gives the service provider (in case of offering a connectivity service) or broadcaster (that leases the service) a consolidated overview of the delivered connection performance.

These tables will then form the basis for assessing the fulfilment of an SLA. Additional performance parameters as Packet delay, Packet Delay Variation and Packet Loss Burst Sizes during the consolidation periods may be added to provide a complete means for following up the SLA.

An example of such a presentation is given below:

Time	ES	SES	UAS	BBE	PLBS	PDV 0.1% (us)	PDV 99.9% (µs)
09:00	12	0	0	325	1	10	157
09:15	0	95	90	0	33	15	180
09:30	3	15	0	263	10	9	171

4. Performance Parameters

4.1 Media Services - Video

4.1.1 Interfaces

SDI: This is a family of standards that define a Serial Digital Interface based on a coaxial cable, intended to be used for transport of uncompressed digital video and audio in a television studio environment.

SD-SDI: SMPTE 259M is a standard to define a 10-bit Serial Digital Interface (SDI). Within the standard the most commonly used bitrate is 270 Mbit/s for carriage of standard definition television.

HD-SDI: SMPTE 292M expands upon SMPTE 259M and SMPTE 344M allowing for bit rates of 1.485 Gbit/s, and 1.485/1.001 Gbit/s. The factor of 1/1.001 is provided to allow SMPTE 292M to support video formats with frame rates of 59.94 Hz, 29.97 Hz, and 23.98 Hz, in order to be upwards compatible with existing NTSC systems. The 1.485 Gbit/s version of the standard supports other frame rates in widespread use, including 60 Hz, 50 Hz, 30 Hz, 25 Hz, and 24 Hz.

3G-SDI: SMPTE 424M is part of family of standards that define a serial digital interface, commonly known as 3G-SDI. It expands upon SMPTE 259M, SMPTE 344M, and SMPTE 292M allowing for bit-rates of 2.970 Gbit/s and 2.970/1.001 Gbit/s over a single-link coaxial cable. These bit-rates are sufficient for 1080p video at 50 or 60 frames per second. The signal formats carried over SMPTE 424M are specified in SMPTE 425M. Guidance on the use of 3G-SDI is given in EBU TR 002.

Other possible interfaces include Dual Link SDI (SMPTE 372M) and Fibre versions

Future versions intended to support UHDTV formats are currently under standardisation at SMPTE including 6G-SDI, 12G-SDI and 24G-SDI.

ASI: Asynchronous Serial Interface is a streaming data format which often carries an MPEG Transport Stream (MPEG-TS). The ASI signal is the final product of video compression, either MPEG-2 or MPEG-4, ready for transmission to a transmitter or microwave system or other device. Sometimes it is also converted to fibre, RF or SMPTE 310 for other types of transmission.

SMPTE ST 2022 is a family of standards to encapsulate video data (uncompressed or compressed in a transport stream) into RTP/UDP/IP streams:

- **SMPTE ST 2022-1** "Forward Error Correction for Real-Time Video/Audio Transport Over IP Networks"
- **SMPTE ST 2022-2** "Unidirectional Transport of Constant Bit Rate MPEG-2 Transport Streams on IP Networks"
- **SMPTE ST 2022-3** "Unidirectional transport of variable bit rate MPEG-2 Transport Streams on IP Networks"
- **SMPTE ST 2022-4** "Unidirectional Transport of Non-Piecewise Constant Variable Bit Rate MPEG-2 Streams on IP Networks"
- **SMPTE ST 2022-5** "Forward Error Correction for High Bit Rate Media Transport over IP Networks"
- **SMPTE ST 2022-6** "High Bit Rate Media Transport over IP Networks" for carriage of SDI over IP.
- **SMPTE ST 2022-7** "Seamless Protection Switching of SMPTE ST 2022 IP Datagrams"

4.1.2 ES, SES definitions

Video: Indicators in the video picture such as freeze frames, black frames, motion jitter, macro blocking, coding artefacts, glitches and break up may be used to define ES and SES, assuming test equipment and techniques are able to reliably detect such errors.

SDI: There is no commonly-agreed definition of ES, SES for SDI services. However the following performance events may be used for this purpose:

- Loss of n consecutive Timing Reference Signals (TRS)
- n Timing Reference Signals (TRS) errors in a one second window
- n Error Detection and Handling (EDH) errors in a one second window for SD-SDI
- n Cyclic Redundancy Codes (CRC) errors in a one second window for HD-SDI

ASI: ETSI TR 101 290 is the authority for ASI signals. Section 5 defines measurements in 3 categories of importance. Section 6.1 defines Severely Errored Seconds and goes on to define Loss of Service as occurring after 10 or more SES. Specific applications may require different or tighter targets. ETSI TR 101 290 sections 5 and 6 may therefore be used by a BMO as a starting point for negotiation.

4.2 Media Services - Audio

4.2.1 Interfaces

AES3 is a standard used for the transport of digital audio signals between professional audio devices. The interface is a point-to-point unidirectional interface primarily designed to carry monophonic or stereophonic signals in a studio environment at 48 kHz sampling frequency and with a resolution of 20 or 24 bit per sample plus 4 bit of metadata. Transmission media include balanced and unbalanced lines and optical fibre.

AES10 (MADI) – also known as MADI (Multichannel Audio Digital Interface), is an Audio Engineering Society (AES) standard that defines the data format and electrical characteristics of an interface that carries multiple channels of digital audio. The MADI standard includes a bit-level description and has features in common with the two-channel format of AES3. It supports serial digital transmission over coaxial cable or fibre-optic lines of 28, 56, or 64 channels; and sampling rates of up to 96 kHz with resolution of up to 24 bit per channel. Like AES3, it is a point-to-point unidirectional interface (one sender and one receiver). The audio data is almost identical to the AES3 payload, though with more channels.

AES67 – The AES67-2013: “AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability” was published in 2013. In the past, a number of networked audio systems were developed to support high-performance media networking, but despite being based on IP and using alike synchronization and transport mechanisms and payload formats, none of these systems were interoperable. The AES addressed this topic by inaugurating the related X192 Task Group in 2010. The goal was to achieve a standard describing interoperability guidelines based on existing technology and well-established standards.

4.2.2 ES, SES definitions

AES3

Block: AES subframe

Errored Block: AES subframe with Parity error

ES: A second with at least 1 (one) Errored Block and not an SES

SES: A second where X% of the blocks are errored

The threshold value “X” will be dependent on the service context. Since an audio service is very susceptible to small errors the threshold value should be low. Typical value is below 1%.

4.3 Transport Performance Parameters and impact on Media Quality

It is important to note that the performance of a Media Service is closely linked with the performance of the underlying Transport layer.

For example, if IP is used to carry a Transport Stream using SMPTE 2022-2 encapsulation, one IP packet can contain up to 7 Transport Stream packets, therefore a loss of 1 IP packet can cause visible picture quality degradation.

This relationship between the Media layer and the Transport layer clearly has an impact on BMOs procuring Media Services from NSPs; but it also has an impact on BMOs procuring Transport Services such as IP or SDH from NSPs and then self-providing the Media layer on top.

4.4 Transport Services - IP

4.4.1 Interfaces

Internet Protocol (IP) does not inherently have a defined physical interface. IP signals are carried at OSI Layer 3 over a variety of OSI Layer 2/1 link transport mechanisms which themselves have specific physical interfaces, as defined in ISO/IEC 7498-1. A mapping is defined for the carriage of IP packets over each link layer type. For Ethernet the encapsulation of IP PDUs is defined in the IEEE 802.3 specification.

4.4.2 Traffic Performance Parameters

Parameter	Standard definition	Description
Packet Loss Ratio (PLR)	Y.1540: IPLR	Packet Loss has directly negative effects to the transported content since it implies loss of media data. Y.1541 Appendix VIII suggests performance objectives ranging from "one hit per 10 days" to "10 hits per day" for different Digital TV related services. Packet Loss may be mitigated to some degree by using repair mechanisms such as Forward Error Correction, FEC. A re-transmission protocol like TCP can recover lost data, but may lead to high delay and delay variations. Therefore, the above objectives, combined with FEC, lead to IP PLR performance objectives of "Table 3 - Provisional IP network QoS class definitions and network performance objectives", that is a PLR of 10e-5.

Parameter	Standard definition	Description
Packet Error Ratio (PER)	Y.1540: IPER	Packet Error Ratio can be seen as a "subset" of PLR, where errored packets actually arrive at the destination. In general, other layers (e.g. Ethernet) will discard errored packets and hence these will show up as lost packets. In Y.1541 the objective for PER is set to a tenth of the value for PLR "to contribute insignificantly to the overall packet loss". It is thus not a significant parameter to specify.
Packet Reordering Ratio (PRR)	Y.1540: IPRR	Due to different routes in an IP network, packets can be delayed so that they arrive after following, newer packets. As long as the buffer at the receiving side is able to reorder, this can be seen as a kind of jitter with no further effect on the transport than buffer delay.
Packet Loss Burst Size (PLBS)	No standard	The distribution of Packet Loss events is of great importance. Even a very low loss IP connection can lead to highly degraded service performance if the losses come in bursts. For example, using a low loss IP link with PLR of 10e-5 and typical FEC matrix sizes, the performance may degrade from the required "one hit per day" to "one hit per 10 minutes". Hence it is very important to control the Packet Loss Burst Size. Since this parameter is not standardized, a suggested objective is that the PLBS should be less than the maximum column FEC size, e.g. 20 in case of MPEG-2 TS streams protected by SMPTE 2022-1 FEC. Packet Loss Bursts of maximum 20 packets should be recoverable by compliant FEC implementations. Smaller PLBS values than 20 are desirable to save on FEC overhead and delay. If PLBS is not controlled, other mechanisms are required, e.g. "hitless 1+1 protection", where data is copied and sent on two separate network paths, and the two packet streams are merged at the egress. This will however require a more complex service and consume much more network bandwidth than plain FEC.
Packet Loss Burst Distance (PLBD)	No standard	Distance in time between packet loss bursts may affect FEC efficiency if it is shorter than the distance in time between two FEC matrices, since one FEC matrix can only handle one packet loss burst. Hence a suitable recommendation would be to have a distance between packet loss bursts that is larger than (L x D x 8) / C where L and D are the dimension of the FEC matrix and C is the bit rate of the service in bits/s. This definition is in line with the rationale for PLBS, however an exact definition for PLBD is yet to be standardised.
Packet Delay (PD)	Y.1540: IPTD	Packet Delay may affect the feasibility of interactive services and remote productions since these services are sensitive to the Round Trip Delay time. The requirements on this parameter are very service-specific.
Packet Delay Variation (PDV)*	Y.1540: IPDV	Packet Delay Variation affects timing-sensitive applications due to the inability to reproduce the service clock to a sufficiently high degree. A jitter buffer, in conjunction with a clock recovery algorithm, is used to recover the service clock, but can only do this to a certain degree. The clock recovery mechanism has a low pass filter characteristics and can hence only filter out short time variations in the service clock, hence long time variations, or wander, is a remaining problem that will depend on the PDV. The relative impact will depend on the type of service. Some services, i.e. studio production or high-end contribution have strict requirements on the wander, while other applications such as IPTV distribution have much lower requirements. Note also that high PDV may lead to Packet Loss if the jitter buffer cannot handle the delay variation.

Parameter	Standard definition	Description
Delay Factor (DF)	IETF RFC 4445	This is the maximum difference, observed at the end of each media stream packet, between the arrival of media data and the drain (i.e. consumption) of media data. It indicates how long a media stream must be buffered at its nominal bit rate to prevent packet loss. This time may also be seen as a measure of the minimum size of the buffer required at the next downstream node to accommodate stream jitter. Display of DF with a resolution of tenths of milliseconds is recommended. The DF value must be updated and displayed at the end of a selected time interval (typically 1 s). The above definition is applicable to CBR streams only. For this reason the EBU has proposed a variant, named Time-Stamped Delay Factor (TS-DF), applicable also to VBR media streams [ref. EBU Tech 3337].

* *The variation in packet delay is sometimes called "jitter". This term, however, causes confusion because it is used in different ways by different groups of people. In this document we refer to the term "Packet Delay Variation" which is more precise, and reserve the terms jitter (and wander) for properties of a signal that is closely related to the timing or clocking of the signal.*

4.4.3 ES, SES definitions

Block: An IP packet.

ES: A second that is not a Severely Errored Second but is subject to one or more lost blocks (i.e. packets)

SES: A second is defined as Severely Errored if the ratio between the number of lost packets (blocks) at the egress of a connection to the number of delivered packets at the ingress of a connection during the second is greater than a defined threshold. The value of the threshold depends on the usage context, thus there is no standardised definition for this.

BBE: A count of lost packets during Errored Seconds

Timing-related parameters (PD/PDV) also affect the performance but these are currently not captured within G.826. A method for calculating ES/SES for timing related performance “events” remains to be formalised. Until then it is recommended to simply report these parameters as-is. For example, Peak-to-peak or 99.9% percentile of PDV during the current 15 minute window.

4.5 Transport Services - Ethernet (Data Link Layer)

4.5.1 Interfaces

Ethernet is a family of computer networking technologies for local area networks (LANs). Ethernet was commercially introduced in 1980 and standardized in 1983 as IEEE 802.3. The standard has subsequently evolved to include many derivatives. The following is a list of the most commonly used:

Copper: 10Base-T, 100Base-TX, 1000Base-T. Generally implemented using Category 5e or Category 6 cabling, as defined in ISO/IEC 11801

Multi-mode fibre: 100Base-SX (up to 300 m), 1000Base-SX, 10GBase-SX (up to 550 m), 40GBase-SR4, 100G-Base-SR10 (up to 150 m). Multi-mode fibre is typically used for indoor cable runs within and between racks of equipment.

Single-mode fibre: 1000Base-LX (up to 10 km) , -XZ (up to 100 km), 10GBase-LR (Up to 10 km), -ER (up to 40 km) or -ZR(Up to 80 km), 40GBase-LR4, 100GBase-LR4 (up to 10 km), 100GBase-ER4 (up to 40 km). Single-mode fibre is typically used for longer indoor cable runs such as within large data centre facilities, as well as for long-range outdoor cable systems.

4.5.2 Performance Parameters

Parameter	Standard definition	Description
Throughput (max fps, byte/s or bit/s)	Y.1563	The total number of Ethernet frames (or octets) transferred during a specified time interval divided by the time interval duration (fps).
Committed Information Rate / Committed Burst Size / Excess Information Rate / Excess Burst Size	MEF10.2	Various bandwidth profile parameters used to define the service quality. The Committed Information Rate defines the average rate in bits/s of Service Frames up to which the network delivers Service Frames and meets the performance objectives defined by the CoS Service Attribute.
Frame Transfer Delay (FTD/FD)	Y.1563/MEF10.2	See "Packet Delay (PD)" above in Network section
Frame Delay Variation (FDV/FDR)	Y.1563/MEF10.2	See "Packet Delay Variation (PDV)" above in Network section
Frame Loss Ratio (FLR/FLR)	Y.1563/MEF10.2	See "Packet Loss Ratio (PLR)" above in Network section
Frame Error Ratio (FER)	Y.1563	See "Packet Error Ratio (PER)" above in Network section
Spurious/Reordered/Duplicate/Replicated Frame events	Y.1563	Second order Frame events, their low probability implies a low practical impact on the media transport

Ethernet frame performance parameters are in all significant aspects equivalent to the corresponding IP packet performance parameters in definitions and impact. Y.1563 defines the parameters but set no performance objectives. MEF 10.2 defines performance parameters and MEF 23.1 specifies performance objectives for different service classes. However, there is no service class defined corresponding to broadcast services.

4.5.3 ES, SES definitions

Block: An Ethernet frame.

ES: A second that is not a Severely Errored Second but is subject to one or more lost blocks (i.e. frames)

SES: A second is defined as Severely Errored if the ratio between the number of lost frames (blocks) at the egress of a connection to the number of delivered frames at the ingress of a connection during the second is greater than a defined threshold. The value of the threshold will depend on the usage context and is for further study.

BBE: A count of lost frames during Errored Seconds

4.6 Transport Services - SDH / SONET (Data Link Layer)

4.6.1 Interfaces

SDH and SONET are transport protocols which operate with a fixed block of data and are used in TDM based networks. “SDH” may be used in the context of a network or a client interface. SONET standards are commonly used in North America and Japan, with SDH in use in other parts of the world. SDH rates have equivalent SONET rates which interwork. SDH is used here as a generic term for both.

An SDH infrastructure has a hierarchy of different fixed speed virtual containers (VC) which support client point to point circuits with dedicated bandwidth through the network.

ITU-T G.707/Y.1322 defines the network interfaces for SDH (the STM-N) interfaces, the frame structure and how client signals are mapped and multiplexed.

ITU-T G.957 defines optical interface parameters for SDH and classifies the “reach” for physical interfaces up to 2.4 Gbit/s (STM-16). The table below from the standard has been extended to show the 155 Mbit/s to 10 Gbit/s STM-64 circuit interface characteristics.

	STM Level	Line Bit Rate (Mbit/s)	Intra-Office	Inter-Office (Short Haul)		Inter-Office (Long Haul)		
Wavelength (nm)			1310	1310	1550	1310	1550	1550
Type of Fibre			G.652	G.652	G.652	G.652	G.652 / G.654	G.652
Distance (km)			≤ 2	~ 15	~ 15	~ 40	~ 80	~ 80
	STM-1	155.52	I-1	S-1.1	S-1.2	L-1.1	L-1.2	L-1.3
	STM-4	622.08	I-4	S-4.1	S-4.2	L-4.1	L-4.2	L-4.3
	STM-16	2488.32	I-16	S-16.1	S-16.2	L-16.1	L-16.2	L-16.3
	STM-64	9953.28	I-64	S-64.1	S-64.2	L-64.1	L-64.2	
	STM-256	39813.12						

4.6.2 Performance Parameters

Parameters	Standard definition	Description
Errored Second Ratio (ESR)	ITU-T G.828	Ratio of ES to total available seconds in a 1 second measurement interval.
Severely Errored Second Ratio (SESR)	ITU-T G.828	Ratio of ES to total available seconds in a 1 second measurement interval.
Background Block Error Ratio (BBER)	ITU-T G.828	Ratio of BBE to total available blocks in a measurement interval.

ITU-T G.828 gives acceptable thresholds for the error rates for ESR, SESR, BBER. The allowed error rate increases with increase of bit rate to maintain a threshold error rate.

The thresholds are used for SLA measurements and trends, so are expected to be measured over a period of 30 days or more for permanent service. For temporary links, measurements have to be set up according the lifecycle of the service.

4.6.3 ES, SES definitions

SDH services supplied via OTN or other optical systems may have "Forward Error Correction" (FEC) active at the optical level for example from an OTN transport system.

The parameters below refer to errors remaining after any FEC correction to the SDH bit stream.

An explanation of the ITU-T model for errors is given in §3.1

ITU-T G.828 specifies limits for errors based on bit rate and a reference model of a 27500 km international circuit.

Hypothetical Reference Path (HRP): A notional international reference circuit of 27500 km, used to define the proportion of errors for the local / national and international portions. The HRP is used to define allowed error thresholds, and how parts of an end to end service contribute to the overall error rates.

SDH digital path: The path or "trail" through an SDH network carrying a VC-x circuit payload and associated overhead.

Block: An SDH frame of consecutive bits. There are 8000 block / second for an SDH service. Bits belonging to a single block may be interspersed with other bits due to multiplexing and overhead. The number of blocks is constant at 155 Mbit/s and above, so the block size increases as the bit rate increases. A block is protected by a CRC to allow errors to be quantified in normal live operation.

EB: An errored block is one with 1 or more bit errors. The CRC cannot distinguish between errors in the payload and errors in overhead bits/ CRC, so all errors are assumed to affect the payload.

ES: An errored second that is not a Severely Errored Second but includes one or more errored blocks.

SES: A severely errored second is one which contains greater than or equal to 30% of errored blocks in 1 second or 1 or more defects.

BBE: An errored block which is not part of a SES.

4.7 Transport Services - Structured WDM and OTN (Physical Layer)

4.7.1 Interfaces

Wavelength Division Multiplexing is a technique to allow a single fibre to carry multiple signals using different colours of light. The colours are normally in the infrared to exploit the frequencies at which optical fibre has the best transmission properties.

In structured WDM, multiple optical signals can be "wrapped" with FEC to improve the error rate using the OTN standards in ITU-T G.709 across a WDM system. The OTN hierarchy is shown below.

OTN Signal	Line Bit Rate (Gbit/s)	Client Server layer	Client Wrapper Bit Rate (Gbit/s)	Typical Client
		ODU0	1.24	1 Gbit/s Ethernet
OTU1	2.67	ODU1	2.50	STM-16 or 2x 1 Gbit/s Ethernet
OTU2	10.71	ODU2	10.04	STM-64
OTU2e	11.09	ODU2e	10.40	10 Gbit/s Ethernet
OTU3	43.02	ODU3	40.32	STM-256 or 4x STM-64
OTU3e2	44.58	ODU3e2	41.79	40 Gbit/s Ethernet or 4x 10 Gbit/s Ethernet
OTU4	112	ODU4	104.79	100 Gbit/s Ethernet

The defined optical signal is OCk (not shown) which is actually transmitted via a fibre or WDM system. OTUk signals include the FEC overhead and carry the corresponding ODU signal.

ITU-T G.709/Y.1331 defines the optical interface parameters for the optical transport hierarchy (OTH) of the OTN. This standard defines how high speed optical channels are presented in order to be carried across an optical network.

4.7.2 Performance Parameters

Parameters	Standard Definition	Description
Errored Second Ratio (ESR)	ITU-T G.828	Ratio of ES to total available seconds in a 1 second measurement interval.
Severely Errored Second Ratio (SESR)	ITU-T G.828	Ratio of SES to total available seconds in a 1 second measurement interval.
Background Block Error Ratio (BBER)	ITU-T G.828	Ratio of BBE to total available blocks in a measurement interval.

ITU-T G.828 gives acceptable thresholds for the error rates for ESR, SESR, BBER. The allowed error rate increases with increase of bit rate to maintain a threshold error rate.

The thresholds are used for SLA measurements and trends, so are expected to be measured over a period of 30 days or more for permanent service. For temporary links, measurements have to be set up according the lifecycle of the service.

4.7.3 ES, SES definitions

Structured WDM systems have FEC active at the optical OTUk level. Error levels are measured after any FEC correction. The parameters below refer to errors remaining after any FEC correction to the client bit stream. Errors are measured independently in each direction. Errors are only measured when the path is available.

The OTN method of measuring errors follows that described in §3.1.

ITU-T G.8201 specifies limits for errors based on bit rate and a reference model of a 27500 km international circuit.

Hypothetical Reference Path (HRP): A notional international reference circuit of 27500 km, used to define the proportion of errors for the local / national and international portions. The HRP is used to define allowed error thresholds, and how parts of an end to end service contribute to the overall error rates.

SDH digital path: The path or “trail” through an SDH network carrying a VC-x circuit payload and associated overhead.

Block: An OTN OPU frame of 4 x 3810 octet (121920 bits). Bits belonging to a single block may be interspersed with other bits due to multiplexing and overhead. A block is protected by a “Bit Interleaved Parity” BIP check to allow errors to be quantified in normal live operation.

EB: An errored block is one with 1 or more bit errors. The BIP only checks the client payload, so header errors are not included.

ES: An errored second that is not a Severely Errored Second but includes one or more errored blocks.

SES: A severely errored second is one which contains greater than or equal to 15% of errored blocks in 1 second or 1 or more defects.

SESR: An error ratio based on SES. The end to end SESR is expected to be better than 0.002 irrespective of client signal bit rate.

BBE: An errored block which is not part of a SES.

BBER: An error ratio based on BBE. The end to end BBER is expected to be better than 2.5 x 10^6 irrespective of client signal bit rate.

4.8 Transport Services - Dark Fibre & bit transparent WDM (Physical Layer)

4.8.1 Interfaces

Dark fibre: Many types of connectors are in use and common types are SC and LC.

Bit transparent WDM: Wavelength Division multiplexing is a technique to allow a single fibre to carry multiple signals using different colours of light. The colours are normally in the infrared to exploit the frequencies at which optical fibre has the best transmission properties.

WDM systems combine and split the light from different interfaces to separate client signals. Practical WDM systems for any distances above a couple of km are deployed using single mode fibre.

The ITU-T defines standard spectral grids for optical frequencies used in Dense WDM (DWDM, in ITU-T G.694.1) and coarse WDM (CWDM, in ITU-T G.694.2).

4.8.2 Performance Parameters

Parameters	Standard Definition	Description
End-to-end Attenuation (dB)	ITU-T G.652-653-655	The attenuation is the reduction in intensity of the light beam with respect to distance travelled through a transmission medium. The provider normally specifies a maximum attenuation at 1310 nm and 1550 nm wavelengths, including connector- and splicing loss at the termination points.

Other performance parameters may be important for specific applications, for example dispersion, polarisation, optical path penalty, etc.

4.8.3 Unavailability definition

Since the fibre is carrying light (analogue signal), the notion of Availability can be determined by its state of working or not and the notion of Error Second is not appropriate. A dark fibre will be Unavailable if it is deteriorated or cut, for instance if the Attenuation falls below a certain threshold (or is infinite if cut).

Service Level Agreements often specifies availability over an extended period, e.g. one year. This is because when a fibre is cut, it may require extensive physical repair work. On the other hand, the risk of many small unavailability periods is close to zero, i.e. the failure pattern is very few down time periods, but longer unavailability periods than for transmission services. In case of fibre cuts, down time can be measured in hours or days, or even weeks in case of sub sea cables.

5. Conclusion

For Broadcast Media Organisations the fundamental business requirement when purchasing services from providers is to get a high quality and reliable service that supports media specific demands.

SLA negotiation should be seen as a communication opportunity to provide a clear and measurable description of the expected service level. A common understanding by both parties can avoid disputes when service does not meet the customer's expectations.

From a business point of view, an SLA is about managing risks and balancing the cost with the correct level of service. When it comes to solving faults and SLA breaches, it is a useful procedural guideline to allow both parties to sort things out when the network service does not behave as stipulated.

This low-level technical document, prepared by the EBU Service Level Agreement (SLA) Project Group as part of the Strategic Programme on Future Networks and Storage Systems (FNS), is intended to complement the high-level document ‘EBU Tech 3361-1 Service Level Agreement for Media Transport Services - High Level Guidelines’, and its glossary ‘EBU Tech 3361-2’.

Feedback on all three documents, and suggestions for amendments or improvements, are welcomed. Please email: info@tech.ebu.ch.

Page intentionally left blank. This document is paginated for two sided printing

Annex A: Path protection architectures

This appendix provides more detail on network path architecture, and explains how this affects the availability of such a network.

Path protection aims to keep a service running in the event of failure of a link or node. Failures can be caused by simple equipment breakdown, or by some kind of damage resulting from fire, flood or the cutting of a cable whether accidentally or maliciously.

Each piece of equipment used in a high-availability service should be designed to have a high 'Mean-Time Between Failure' (MTBF). To achieve this, equipment manufacturers use internal protection techniques including multiple power supplies, multiple internal signal busses, and modular backplane and card structures. When procuring a managed service with a stated availability, it is useful for a BMO to understand the MTBF of the key items of equipment proposed by the NSP as this may provide an indication of the frequency of failures experienced in the service. However, it is the way these items of equipment are connected to form a transmission link, together with the time taken for the NSP to repair a failed piece of equipment, which determines the overall service availability, so it is more important for the BMO to assess the high-level service architecture than the internal architecture of each device.

Path protection is created using a combination of duplicated equipment and mechanisms for switching traffic in the event of equipment failure. At its most basic, path protection can be seen as providing more than one path for the traffic between entry and egress points. NSPs use the expressions '1+0' and '1+1' to differentiate between unprotected and protected paths. Four example protection architectures are shown in the diagrams at the end of this section.

However, there are many complexities underlying this apparently simple definition. Questions a BMO needs to ask when assessing a protected service offering include:

- Is the redundant path engineered with the same performance as the main path, including bitrate, latency, jitter and QoS traffic priorities?
- Is my traffic 'live-live' on both 'A' and 'B' paths, or just 'live-standby'? Will the 'Standby' be available when I need it?
- Does the service provide dual handoffs ('A' and 'B') at both ends, or only a single handoff with path protection inside the service? Can I use the 'A' and 'B' paths to carry different traffic if I wish?
- Is the entire path protected, including each intermediate node, or just the cables/microwave links?
- What layer in the service is used to provide the protection switching? Examples include Optical, SDH, Ethernet, IP, ASI.
- If the switching is in a packet-based layer, is switching performed per stream or per packet?
- Is the switchover between main and standby paths automatic? What is the switchover time, and will this cause a glitch in the service? Does the service automatically switch back to the main path when the fault is fixed, and if so can this create a 'flip-flop' situation in the event of an intermittent fault? If automatic switching is deployed, how is a failure detected?
- Is the redundant path route sufficiently separated from the main route such that an incident on the main path cannot also take down the redundant path – e.g. fire, flood, power failure? Note this feature differentiates a service with 'diversity' from a service with 'separacy'.

Typically protection is engineered to survive one failure at a time, as the probability of multiple simultaneous failures is low. However for long paths passing through multiple links and nodes, it is wise to engineer the capability for fail-over on a per-segment basis so that failure of the 'A' path in one segment coupled with failure of the 'B' path in another segment does not cause total service failure. The MSO should assess where in such a long path switchover is allowed to occur, and the risk of convoluted paths creating unexpected failure scenarios.

For very high availability services, it is common for a NSP to provide double-protection, using protection mechanisms in two different layers of the solution. An example of this is where the 'A' and 'B' path are switched in the IP or ASI domain, but each runs over different SDH line systems which are themselves equipped with SDH ring-resilience switching. This effectively creates a (1+1)+(1+1) architecture, but the MSO must be absolutely certain in this situation what possible combination of paths and nodes are allowed for the service by the NSP, as such a complex mesh of paths can create scenarios where all traffic is traversing one node.

Example protection architectures:

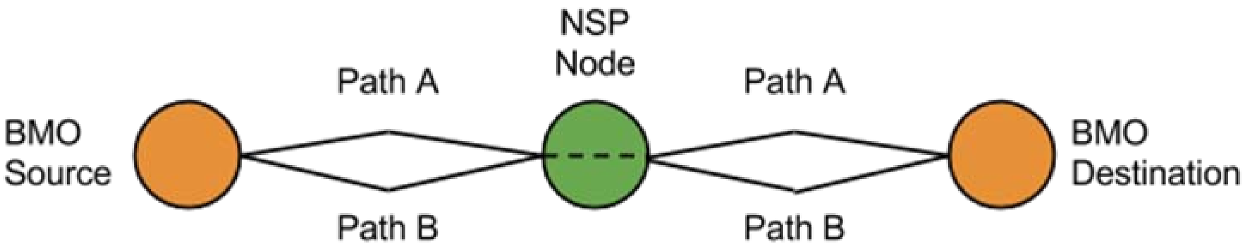

Figure A1: Single handoffs, protected links, single node with single traversal

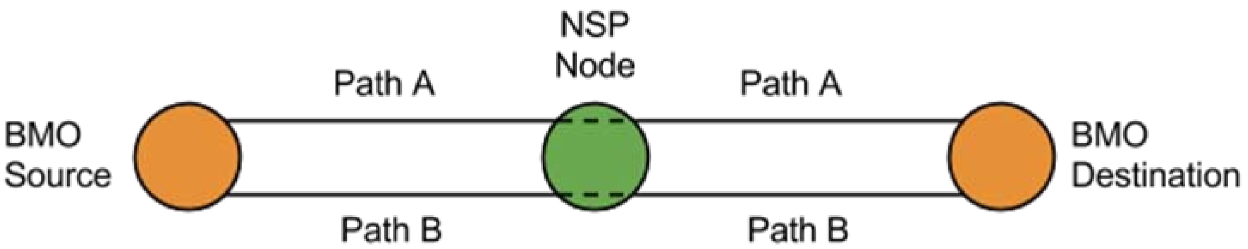

Figure A2: Dual handoffs, protected links, single node with dual traversal

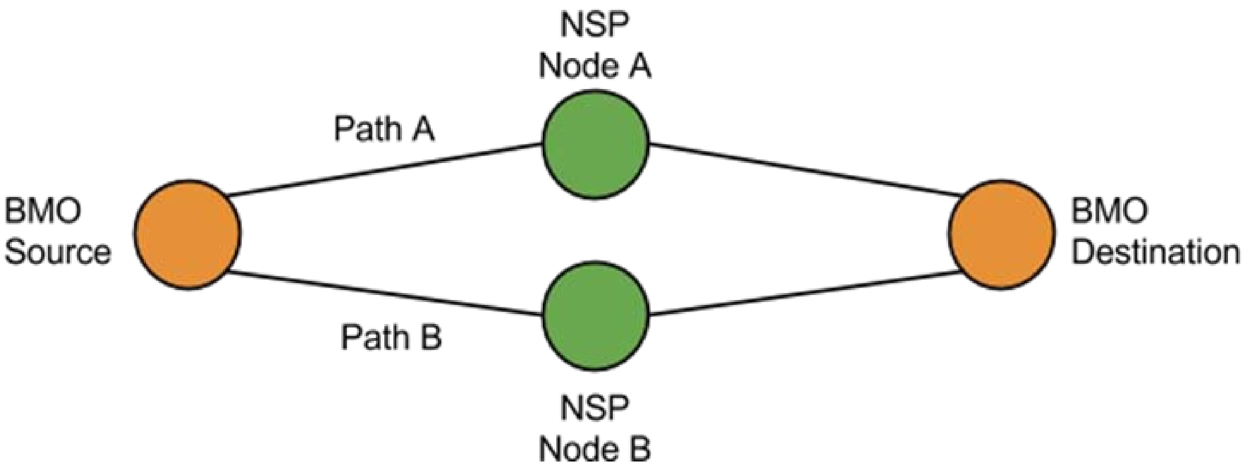

Figure A3: Dual handoffs, dual links, dual nodes

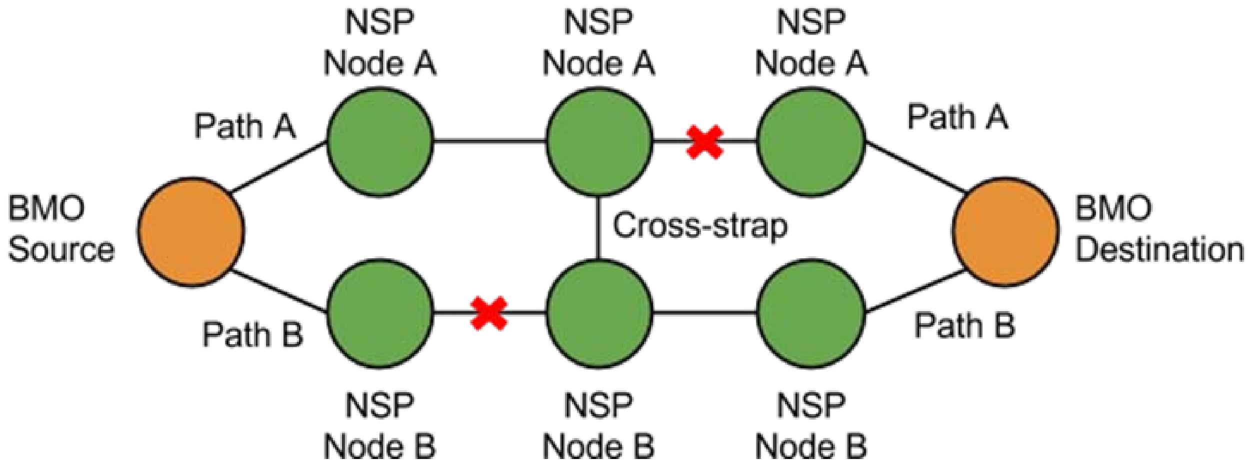

Figure A4: Dual handoffs, dual links, dual nodes, multiple segments with cross-strap allowing service restoration in case of multiple failures (red crosses showing example failure scenario)