The scaling required to normalize the average power is dependent only on the size of the constellation. It is represented by $\chi(b_i)$ and is specified in clause 10.3.4.1.

The gain adjuster g_i is used to equalize the SNR margin over the subcarriers in use and is specified in clause 10.3.4.2.

The PSD shaping mechanism is based on tss_i coefficients and is specified in clause 10.3.4.3. The shaping by a tss_i value is in addition to any shaping introduced by time-domain filters (if used).

For subcarriers in the MEDLEY set, each constellation point (X_i, Y_i), corresponding to the complex value $X_i + jY_i$ at the output of the constellation mapper, shall be scaled by the power-normalization factor $\chi(b_i)$, the gain adjuster g_i, and a frequency-domain spectrum shaping coefficient tss_i to result in a complex number Z_i, defined as:

$$Z_i = g_i \times tss_i \times \chi(b_i) \times (X_i + jY_i)$$

10.3.4.1 Power normalization

The values (X, Y) shall be scaled such that all constellations, regardless of size, have the same average power. The required scaling, $\chi(b_i)$, is a function only of the constellation size.

10.3.4.2 The gain adjuster

The gain g_i is intended for fine gain adjustment within a range from approximately 0.1888 to 1.33, which may be used to equalize the SNR margin for all subcarriers. The g_i values in dB shall be defined as the $20\times\log_{10}(g_i)$, thus g_i values of 0.1888 and 1.33 in linear scale correspond to g_i values of −14.5 dB and of +2.5 dB, respectively. The values of g_i for all MEDLEY subcarriers shall be assigned during initialization, as described in clause 12.3.5 and stored in the bits-and-gains table specified in clause 10.3.1 (b_i and g_i values). The g_i values may also be updated during showtime via an OLR procedure described in clauses 13 and 11.2.3.3.

The g_i settings (in the bits-and-gains table) shall comply with the following requirements:

- If $b_i > 0$, then g_i shall be in the [−14.5 to +2.5] (dB) range;
- If $b_i > 0$, then the linear average of the g_i^2's in any band (as specified during the ITU-T G.994.1 handshake phase of initialization, see clause 12.3.2) shall be ≤ 1;
- If $b_i = 0$, then g_i shall be equal to 0 (linear) or in the [−14.5 to 0] (dB) range;
- The gain adjustments shall be set in accordance with service priorities specified in clause 12.3.7.

For subcarriers not in the MEDLEY set, see Table 10-4.

10.3.4.2.1 Nominal aggregate transmit power (NOMATP)

The nominal aggregate transmit power (NOMATP) shall be computed by the following equation:

$$\text{NOMATP} = 10\log_{10} \Delta f + 10\log_{10}\left(\sum_{i\in\text{MEDLEY set}}\left(10^{\frac{\text{MREFPSD}[i]}{10}} g_i^2\right)\right)$$

where MREFPSD[i] and g_i are, respectively, the values of MREFPSD in dBm/Hz and gain (linear scale) for subcarrier i from the MEDLEY set (see clause 12.3.3.2.1.3), and Δf is the subcarrier spacing in Hz.

The downstream NOMATP (NOMATPds) shall be computed for subcarriers from the downstream MEDLEY set (MEDLEYds). The upstream NOMATP (NOMATPus) shall be computed for subcarriers from the upstream MEDLEY set (MEDLEYus).

The downstream maximum nominal aggregate transmit power during initialization and showtime (parameter MAXNOMATPds) is defined by the CO-MIB as specified in [ITU-T G.997.1]. The

MAXNOMATPds settings in the CO-MIB shall not exceed the maximum downstream aggregate transmit power specified in Table 6-1.

The upstream maximum nominal aggregate transmit power during initialization and showtime is defined by the control parameter *MAXNOMATPus*. The control parameter *MAXNOMATPus* is determined by the maximum aggregate upstream transmit power specified in Table 6-1.

NOTE – There is no configuration parameter MAXNOMATPus defined in the CO-MIB (see Table 7-15 of [ITU-T G.997.1]).

The g_i settings at the VTU-O and VTU-R shall be such that the values of NOMATPds and NOMATPus do not exceed, respectively, the CO-MIB parameter MAXNOMATPds and the control parameter *MAXNOMATPus*. To assist the proper gain setting at the VTU-O, the MAXNOMATPds is communicated from the VTU-O to the VTU-R during the channel discovery phase.

10.3.4.3 Frequency-domain transmit spectrum shaping (tss_i)

The tss_i are intended for frequency-domain spectrum shaping, both upstream and downstream. The tss_i values are vendor discretionary and shall be in the range between 0 and 1 (linear) in steps of $\frac{1}{1024}$. The tss_i values shall be set such that the highest tss_i value across all subcarriers is 1. Smaller values of tss_i provide attenuation, and the value $tss_i = 0$ corresponds to no power transmitted on the particular subcarrier. If no frequency-domain spectrum shaping is applied, the tss_i values shall be equal to 1 for all subcarriers.

The tss_i values in dB (log_tss_i) are defined as $20 \times \log_{10}(tss_i)$ and shall be converted to linear values of tss_i using the equation:

$$tss_i = \frac{\text{Round}\left(1024 \times 10^{\frac{log\_tss_i}{20}}\right)}{1024}$$

The values of tss_i for the given direction of transmission shall be determined by the transmitting VTU, and shall be defined as a set of breakpoints $\{(i_1, log\_tss_{i1}) ..., (i_n, log\_tss_{in})\}$, where i is the subcarrier index. This set shall be communicated to the receiving VTU during the channel discovery phase of the initialization using O-PRM and R-PRM messages, as described in clause 12.3.3.2. Both transmitting and receiving VTUs shall derive the tss_i values for subcarriers between the breakpoints using linear interpolation of the defined log_tss_i values over the linear scale of subcarrier indexes. The receiving VTU shall assign tss_i values equal to tss_{in} for $i > i_n$, and equal to tss_{i1} for $i < i_1$.

The obtained values of tss_i are relevant only for subcarriers that are actually transmitted. The receiver shall ignore the tss_i values that are either received or obtained by interpolation for the subcarriers that are not used for transmission (Z_i=0, see Table 10-4).

The combined accuracy of the linear interpolation of log_tss_i values and of the conversion to linear tss_i values shall be less than one half LSB for the 10-bit representation format of the linear tss_i values. No error shall be introduced when log_tss_i equals 0 dB or is interpolated between log_tss_i values that equal 0 dB.

The transmitters of the VTU-O and VTU-R, respectively, shall set the tss_i values such that, prior to the gain adjustment (i.e., assuming g_i =1), the PSD of the transmit signal as measured in the reference impedance at the U interface, from the start of the training phase and for the remainder of initialization, shall not deviate from the values of MREFPSDds and MREFPSDus, communicated in O-PRM and R-PRM, respectively, by more than 1 dB (parameter "MEDLEY reference PSD", see clause 12.3.3.2). Thus, tss_i settings shall take into consideration any additional spectrum shaping caused by time-domain filters and analog filters included in the transmission path between the output of the modulator and U interface.

10.3.4.4 Summary of subcarrier constellation mapping and constellation point scaling

Table 10-4 summarizes the subcarrier constellation mapping and constellation point scaling requirements for the stages of initialization and during showtime.

Table 10-4 – Summary of subcarrier modulation in initialization and showtime

Phase	Subcarrier index (i)		Z_i
Channel discovery (12.3.3)	$i \in$ SUPPORTEDCARRIERS		$tss_i \times (X_i + jY_i)$
	$i \notin$ SUPPORTEDCARRIERS		0
Training (12.3.4)	$i \in$ MEDLEY		$tss_i \times (X_i + jY_i)$
	$i \notin$ MEDLEY (Note 1)		0
Channel analysis and exchange (12.3.5)	$i \in$ MEDLEY		$tss_i \times (X_i + jY_i)$
	$i \notin$ MEDLEY		0
Showtime	$i \in$ MEDLEY	$b_i > 0$, $g_i > 0$	$g_i \times tss_i \times \chi(b_i) \times (X_i + jY_i)$
		Monitored subcarriers ($b_i = 0$, $g_i > 0$, modulated by 4-QAM)	$g_i \times tss_i \times \chi(b = 2) \times (X_i + Y_i)$
		Pilot tones ($b_i = 0$, $g_i > 0$, modulated by 4-QAM)	$g_i \times tss_i \times \chi(b = 2) \times (X_i + Y_i)$
		Others with $b_i = 0$, $g_i = 0$	0
	$i \notin$ MEDLEY	$i \in$ SUPPORTEDCARRIERS, and $i \in$ BLACKOUT	0
		$i \in$ SUPPORTEDCARRIERS, and $i \notin$ BLACKOUT	Vendor discretionary (Note 2)
		$i \notin$ SUPPORTEDCARRIERS	0

NOTE 1 – The O-P-TEQ and R-P-TEQ signals used during the training phase include subcarriers that are outside the MEDLEY set. See clause 12.3.4.3 for details.

NOTE 2 – The PSD of vendor-discretionary signals on these subcarriers shall be below MREFMASK by 10 dB.

10.4 Modulation

10.4.1 Data subcarriers

The subcarriers shall be indexed from $i = 0$ to $i = MSI$, where MSI is the index of the highest loaded subcarrier (i.e., the maximum index in the MEDLEY set). The values of MSI may be different for upstream and downstream transmission and are denoted as MSI_{us} and MSI_{ds} respectively. The index of the highest loaded subcarrier (MSI_{us} or MSI_{ds}) will be restricted by the selected profile and band plan as shown in Table 6-1. Specifically, MSI_{us} shall be equal to or lower than the "index of the highest supported upstream data-bearing subcarrier" (6.2.10) and MSI_{ds} shall be equal to or lower than the "index of the highest supported downstream data-bearing subcarrier" (6.2.9). Transmission will take place on NSC subcarriers, with $NSC_{us} \leq MSI_{us}$ and $NSC_{ds} \leq MSI_{ds}$; the subcarrier with index i=0 shall not be used. $NSC_{us} + NSC_{ds}$ shall always be less than 4096.

The subcarriers to be used for data transport in the upstream and downstream directions (MEDLEYus and MEDLEYds sets, respectively) shall be determined during initialization, as specified in clause 12.3.3.

NOTE – The subcarriers used for data transmission depend on channel characteristics, such as loop attenuation and noise, and on the specific requirements on the PSD of the transmit signal, such as notching of amateur radio bands, PSD reduction at low frequencies to share the loop with POTS or ISDN, and others.

10.4.2 Subcarrier spacing

Subcarrier spacing is the frequency spacing, Δf, between the subcarriers. The subcarriers shall be centered at frequencies $f = i \times \Delta f$. The subcarrier index i takes the values i = 0, 1, 2, ..., *MSI*. Valid values of subcarrier spacing are 4.3125 kHz and 8.625 kHz, both with a tolerance of ±50 ppm.

Subcarrier spacing is profile dependent (see Table 6-1).

10.4.3 Modulation by the inverse discrete Fourier transform (IDFT)

The IDFT is used to modulate the output of the symbol encoder onto the DMT subcarriers. It converts the *NSC* complex values Z_i (as defined in clause 10.3.4) generated by the symbol encoder (frequency domain representation) into $2N$ real values x_n ($n = 0, 1, ..., 2N - 1$), which is a time domain representation. The conversion shall be performed with a $2N$ point IDFT, with $N - 1 \geq MSI$, as:

$$x_n = \sum_{i=0}^{2N-1} \exp\left(j \cdot 2 \cdot \pi \cdot \frac{n \cdot i}{2 \cdot N} \right) \cdot Z_i \qquad \text{for } n = 0 \text{ to } 2N - 1$$

The valid values of N are $N = 2^{n+5}$, where n can take integer values from 0 to 7. The values of N used for upstream and downstream are exchanged during initialization (see clauses 12.3.2, 12.3.3.2.1.3, and 12.3.3.2.2.3).

For subcarrier indices i that are not in the MEDLEY set and for $MSI < i < N$, the corresponding values of Z_i are not generated by the symbol encoder. These values are vendor discretionary, but shall comply with the constraints given in Table 10-4. Z_0 shall always be equal to zero and Z_N shall always be a real value.

In order to generate real values of x_n, the input values Z_i, where $i = 0, 1, ..., N - 1$ and $Z_0 = 0$, shall be further augmented so that the vector Z_i has a Hermitian symmetry:

$$Z_i = \text{conj}\,(Z_{2N-i}) \qquad \text{for } i = N + 1 \text{ to } 2N - 1$$

NOTE – Different values of N result in different transmit signal images above the Nyquist frequency. Knowledge of how the additional Z_i values are defined allows the receiver to better estimate the channel during initialization.

10.4.4 Cyclic extension and windowing

The transmit DMT symbol shall be constructed from the IDFT samples x_n using the following rules.

The last L_{CP} samples of the IDFT output x_n shall be prepended to the $2N$ output IDFT samples x_n as the cyclic prefix (CP). The first L_{CS} samples of x_n shall be appended to the block of $x_n + L_{CP}$ samples as the cyclic suffix (CS). The first β samples of the cyclic prefix and last β samples of the cyclic suffix shall be used for shaping the envelope of the transmitted signal (windowing). The values of the window samples are vendor discretionary. The maximum value of β shall be min(N/16, 255). The windowed parts (β samples) of consecutive symbols shall overlap and be added to one another.

Figure 10-14 summarizes all of the operations that shall be performed by the transmitter to construct the DMT symbol.

The cyclic extension (CE) length is defined as $L_{CE} = L_{CP} + L_{CS} - \beta$. The values L_{CP}, L_{CS} and β shall be set in order to satisfy the equation $L_{CE} = (L_{CP} + L_{CS} - \beta) = m \times N/32$, where valid values of m are integers between 2 and 16, inclusive. Support for the value of m = 5 is mandatory. In all cases, the following relations shall hold: $\beta < L_{CP}$ and $\beta < L_{CS}$.

NOTE – Partitioning between the CS and CP is vendor discretionary. The specific settings of the CE and CP are exchanged during initialization.

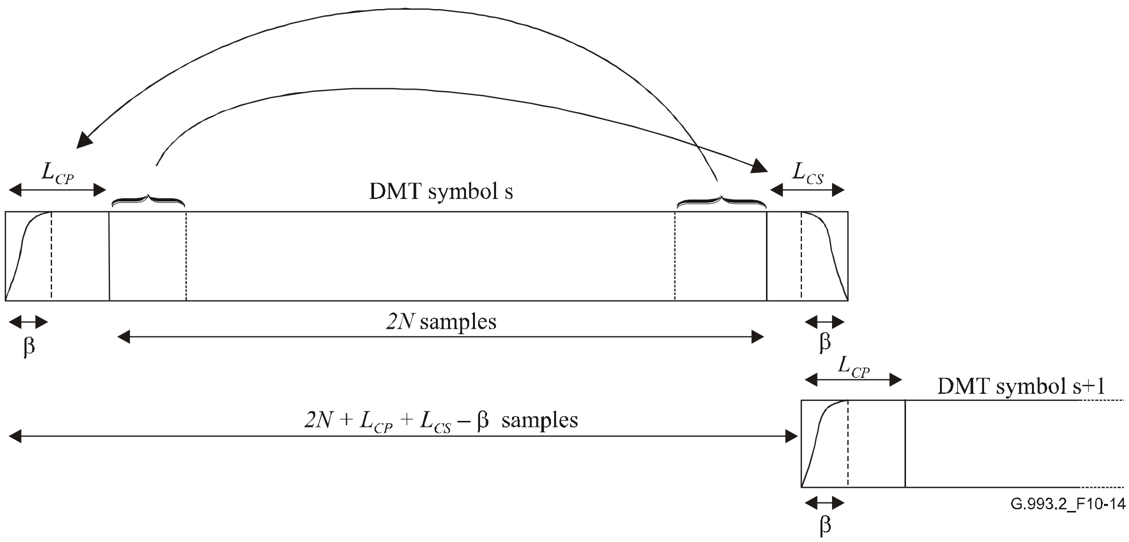

Figure 10-14 – Cyclic extension, windowing and overlap of DMT symbols

For a given setting of the CE length and window length β, the DMT symbols will be transmitted at a symbol rate equal to:

$$f_{DMT} = \frac{2N \times \Delta f}{2N + L_{CP} + L_{CS} - \beta} = \frac{2N \times \Delta f}{2N + L_{CE}}$$

If the CE length corresponds to $m = 5$, this results in symbol rates of 4 ksymbols/s for $\Delta f = 4.3125$ kHz and 8 ksymbols/s for $\Delta f = 8.625$ kHz, independent of the sampling rate used.

The data symbol rate is equal to:

$$f_s = \frac{2N \times \Delta f}{2N + L_{CP} + L_{CS} - \beta} \times \frac{256}{257}$$

The equivalent 4k DMT symbol length (denoted as T_{4k}) is defined as:

$$T_{4k} = \frac{1}{f_{DMT}}$$ for Δf = 4.3125 kHz, and

$$T_{4k} = \frac{2}{f_{DMT}}$$ for Δf = 8.625 kHz.

If the CE length corresponds to $m = 5$, this results in an equivalent 4k DMT symbol length of 250 μs, independent of the subcarrier spacing and sampling rate used.

10.4.5 Synchronization

10.4.5.1 Pilot tones

The VTU-R may select one or more subcarriers to use for timing recovery, called "pilot tones". Pilot tones are selected separately for initialization and showtime.

Pilot tones during initialization: The VTU-R may select initialization pilot tones by indicating its selection of pilot tones in R-MSG 1 (see clause 12.3.3.2.2.1). Initialization pilot tones are used for initialization signals O-P-PILOT1, O-P-PILOT2, O-P-PILOT3 and O-P-ECT as specified in clauses 12.3.3 and 12.3.4. The total number of initialization pilot tones shall not exceed 16.

Pilot tones during showtime: The VTU-R may select showtime pilot tones by indicating its selection of pilot tones in R-PMD (see clause 12.3.5.2.2.4). The VTU-O shall transmit on the selected subcarriers the value of 00 using 4-QAM modulation during every data symbol of showtime. The constellation point scaling for the pilot tone(s) shall follow the same rules as for data carrying subcarriers described in clause 10.3.4. The total number of pilot tones shall not exceed 16. Pilot tones are not transmitted on sync symbols (see clause 10.5.1).

10.4.5.2 VTU-R timing

The VTU-R shall perform loop timing (see clause 3.30).

10.4.5.3 Timing advance

The VTU-R shall be capable of implementing a timing offset between transmit and receive DMT symbols, called timing advance (TA). It shall set the TX symbol TA reference point prior to the RX symbol TA reference point by the value of TA, see Figure 10-15. For the purpose of implementing TA, the TX symbol TA reference point is floor(β_{us}/2) samples after the first sample of the cyclic prefix. Similarly, for the purpose of implementing TA, the RX symbol TA reference point is floor(β_{ds}/2) samples after the estimated first sample of the cyclic prefix. The estimation of the first sample of the received symbol is vendor discretionary and may depend on loop conditions. However, the VTU-R should make its best effort to meet the TA at the U interface. The TA shall be calculated and set during initialization, as specified in clauses 12.3.3 and 12.3.4.

If the value of TA is exactly equal to the propagation delay from the VTU-O to the VTU-R, it will force the VTU-O and VTU-R to start transmission of DMT symbols in opposite directions simultaneously (i.e., the DMT symbols in the downstream and upstream transmission directions start at the same absolute time). This results in orthogonality between transmitted and received DMT symbols when the minimum value of CE length is used.

NOTE – To obtain the desired orthogonality between transmit and receive signals with the minimum value of CE length, the value of TA should apply at the U interface.

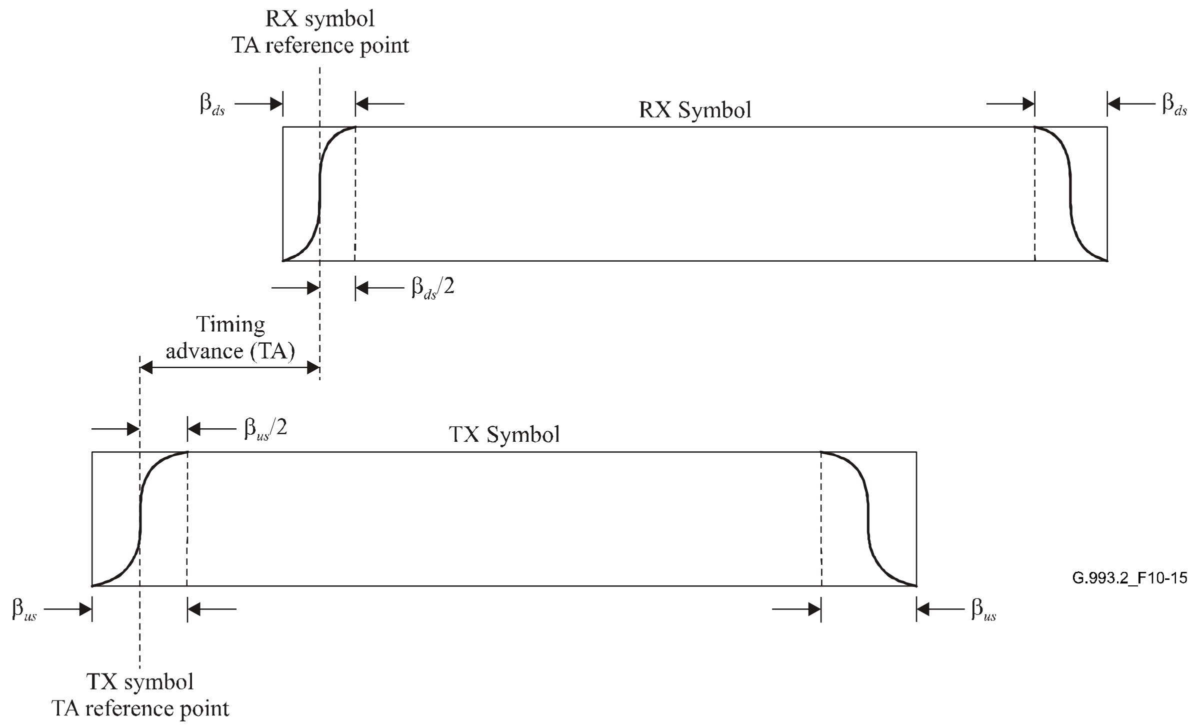

Figure 10-15 – Illustration of TA (VTU-R)

10.4.5.4 Synchronous mode

Support of synchronous mode is optional. In synchronous mode, the out-of-band near-end crosstalk (NEXT) generated by the VDSL2 systems operating in synchronous mode will be nearly orthogonal to the VDSL2 signals received by other VTUs operating in synchronous mode. Therefore, the NEXT will not significantly degrade the SNR on other lines in synchronous mode.

In synchronous mode, all VTU-Os shall use the same subcarrier spacing and symbol rate, and shall start transmission of DMT symbols at the same time on all of the lines in the synchronized group. The transmit symbol clocks shall be phase-synchronous at all VTU-Os with a 1 µs maximum phase error tolerance.

In synchronous mode, all VTUs shall use the same value of CE length (see clause 10.4.4). The CE length used for all lines in the synchronized group should have values appropriate for the line in the group that has the largest propagation delay.

10.5 Symbol encoder for sync symbol

10.5.1 Constellation mapper for sync symbol

Each MEDLEY subcarrier of the sync symbol in either transmission direction (MEDLEYds or MEDLEYus; see clauses 12.3.3.2.1.3 and 12.3.3.2.2.3) shall be modulated by two bits from the sync frame (which will be either 00 or 11 for all MEDLEY subcarriers) using the 4-QAM constellation defined in clause 10.3.3.2.1. The constellation points on these subcarriers shall then be rotated based on the 2-bit number provided by the quadrant scrambler described in clause 12.3.6.2. The scrambler shall be used in reset mode (see clause 12.3.6.2.1).

Pilot tones (as specified in clause 10.4.5.1) are not transmitted on sync symbols.

NOTE – The first five and last five symbols of initialization signals O-P-SYNCHRO and R-P-SYNCHRO are identical to a sync symbol modulating a sync frame of all ONEs. The middle five symbols of O-P-SYNCHRO and R-P-SYNCHRO are identical to a sync symbol modulating a sync frame of all ZEROS.

An inversion of the bits in the sync frame (i.e., from all ONES to all ZEROS and vice versa) shall be used to signal on-line reconfiguration timing during showtime, as described in clause 10.5.3.

For the subcarriers in the SUPPORTEDCARRIERS set that are not in the MEDLEY set and are not in the BLACKOUT set, the constellation mapper may select a vendor-discretionary (X, Y) point, which may also change from one sync symbol to another (see Table 10-4).

10.5.2 Constellation point scaling for sync symbol

The $\chi(b_i)$, g_i and tss_i values shall be applied to the sync symbol in the same way as they are applied to data symbols in showtime (see clause 10.3.4).

10.5.3 On-line reconfiguration

The transmitter inserts a sync symbol every 257 symbols, as defined in clause 10.2. Therefore, a sync symbol shall be transmitted after every 256 data symbols.

To signal on-line reconfiguration timing (see clause 13.3), the responding VTU shall send a Syncflag (see clause 3.64).

After the transmission of a Syncflag, the sync frame modulated onto subsequent sync symbols shall remain the same (i.e., either all ONES or all ZEROS) until timing for the next on-line reconfiguration needs to be signalled.

At the beginning of showtime, the first sync symbol transmitted shall be modulated by a sync frame of all ones.

10.6 Symbol encoder for initialization

Encoding of DMT symbols transmitted during the different phases of initialization is specified in clauses 12.3.3.3, 12.3.4.3, and 12.3.5.3. The values of *X* and *Y* of the 4-QAM constellation points during initialization shall be as shown in the constellation diagram of Figure 10-9. These values shall be scaled such that at the output of the constellation mapper the constellation represents the root mean square (rms) energy of a subcarrier transmitted at the relevant PSD level. The applicable PSD levels are specified in clauses 12.3.3.3, 12.3.4.3 and 12.3.5.3.

11 Operation and maintenance (OAM)

11.1 OAM functional model

The operations, administration and maintenance (OAM) reference model of a VDSL2 link, as shown in Figure 11-1, contains OAM entities intended to manage the following transmission entities:

- *VDSL2 Line entity*: The physical transmission entity, which includes the PMD and PMS-TC sublayers;
- *VDSL2 Path entity*: The transport protocol path, which includes the TPS-TC sublayer; and
- *VDSL2 System entity*: The application path, which includes all relevant layers above the TPS-TC.

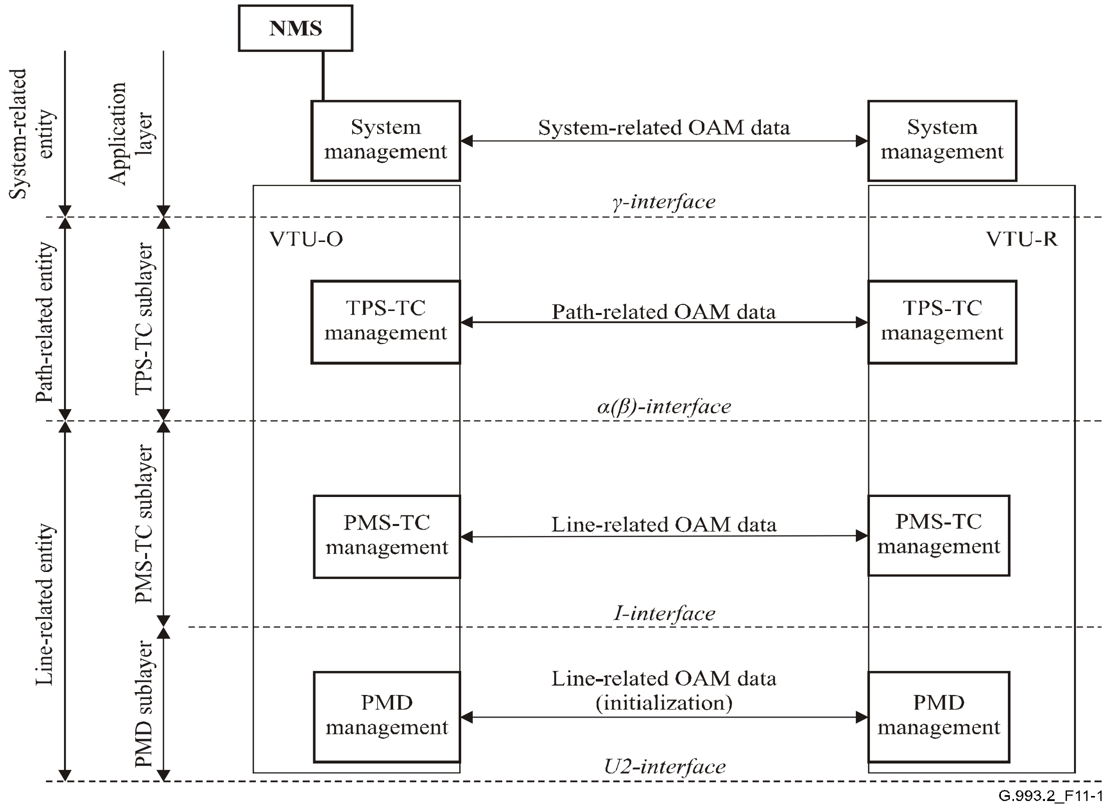

Figure 11-1 – OAM reference model

The peer OAM entities at the VTU-O and VTU-R exchange management data over OAM-dedicated communication channels arranged over the mentioned transmission entities. The NMS, located at the VTU-O, controls the OAM entities at both VTUs, and collects management data from all OAM entities. The OAM flows across the communication channels convey path-related and line-related primitives and parameters, configuration setups, and maintenance commands and acknowledgments.

The functional model of the OAM operation and communication over the VDSL2 link is presented in Figure 11-2. The external OAM interface adapter (EIA) provides the interface to the NMS (Q interface), and the interface with the MIB. The MIB contains all of the management information related to the VDSL2 link. It may be implemented to serve an individual VDSL2 line or to be shared between several lines.

The VME collects the OAM data from and delivers it to all of the VTU transmission entities, thus providing all internal OAM functions for the VTU. It also supports all interactive management functions between the VTU-O and the VTU-R using two OAM-dedicated communication channels:

- Indicator bits (IB) channel; and
- Embedded operations channel (eoc).

The VME interfaces with the EIA, thus exchanging management data with the MIB. The VME functionality is specified in clause 11.2.1. The EIA functions concerning operation with the external interfaces (Ext_OAM_R interface, Q interface), with the MIB, and the interface between VME and EIA are beyond the scope of this Recommendation.

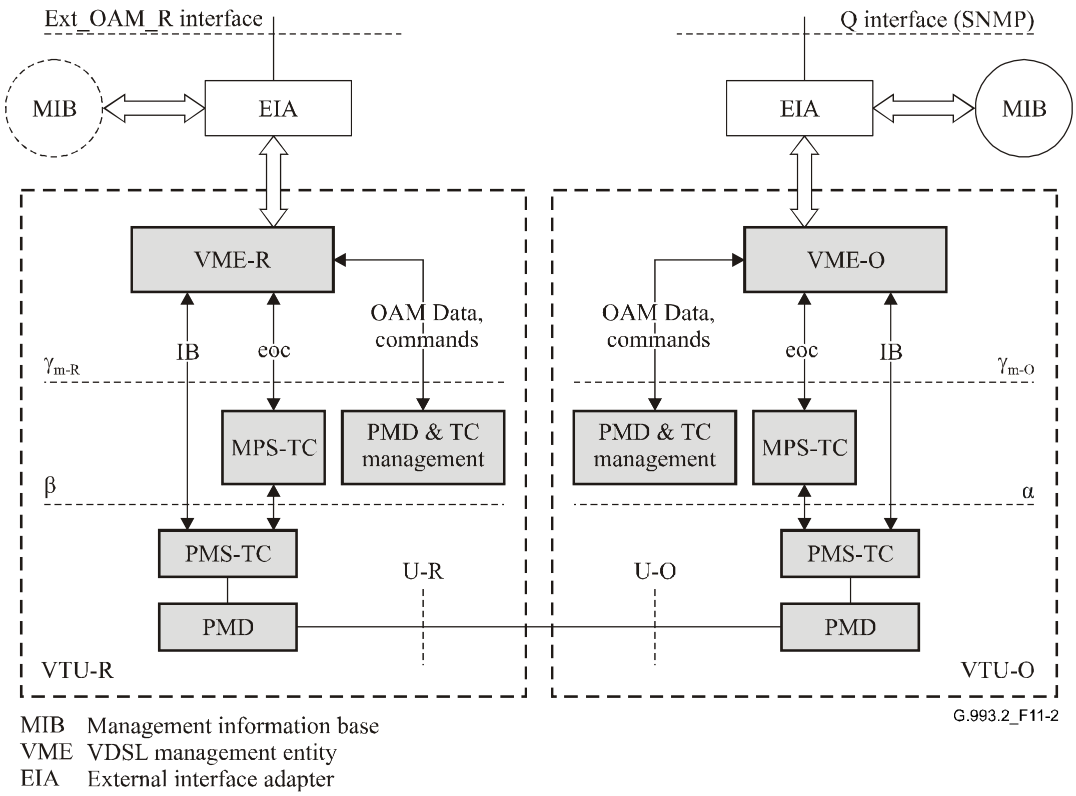

Figure 11-2 – Functional model of OAM of the VDSL2 link

To communicate management data, the VME uses eoc messages (specified in clause 11.2.3) and IB (specified in clause 11.2.4). The eoc messages and IB form a complete set of management data exchanged between the VTU-O and VTU-R, which includes the management data from all data-transmission sublayers of the VTU and the management data incoming from the EIA, including messages sent to the VTU-R. The latter are referred to in [ITU-T G.997.1] and in clause 11.2.3 as a "clear eoc". The interfaces between the VME and the TC sublayer for both OAM communication channels are functional and are defined in clause 8.2.2 (MPS-TC) and in clause 9.5.2.2 (IB). The eoc communication protocol is defined in clause 11.2.2.

The VME sends eoc messages via the γ_m interface to the management TPS-TC (MPS-TC) defined in clause 8.2.2. The MPS-TC encapsulates eoc messages into HDLC frames to transfer them over the VDSL2 link using the PMS-TC overhead channel (the MSG field of the OH frame specified in clause 9.5.2.2). At the receive side, the MPS-TC extracts the received eoc messages from HDLC frames and submits them to the VME via the γ_m interface.

For the IB transport no TPS-TC is needed; the IB are directly mapped to the IB field of the OH frame as specified in clause 9.5.2.2.

11.1.1 OAM communication channels

11.1.1.1 IB channel

The IB channel is shared for communication between the peer OAM entities of the PMD, PMS-TC and TPS-TC. It is intended to transfer time-sensitive primitives (those requiring an immediate action) from the far end. The IB channel operates in a unidirectional mode, i.e., the upstream and downstream directions of the IB channel operate independently, and there are no acknowledgements or retransmissions in the protocol. The IB are specified in clause 11.2.4.

11.1.1.2 eoc

The eoc is shared for communication between the peer OAM entities of the PMD, PMS-TC, TPS-TC and VME (system-related OAM data, such as power-related primitives). The eoc is mostly intended to exchange management data that is not time critical. It is used to transport clear eoc messages and MIB elements specified in [ITU-T G.997.1], to set and query parameters, and to invoke management procedures at the far-end VTU. The eoc provides exchange of the PMD, PMS-TC, TPS-TC and system-related primitives, performance parameters, test parameters, configuration parameters and maintenance commands. The eoc communication protocol is specified in clause 11.2.2.

11.2 VDSL2 management entity (VME)

11.2.1 VME functionality

The VME provides all necessary management functions specified in [ITU-T G.997.1] to communicate with the MIB and with the NMS via the EIA. It shall also manage the OAM communication channels, and support all internal management functions of the VTU, including:

- performance monitoring;
- performance management;
- configuration management; and
- fault management.

The VME shall provide all of the functionality to communicate the management data between the VTU-O and the VTU-R. Specifically, the VME shall:

- originate eoc messages and IB to communicate management data;
- assign priority levels for eoc messages to share the overhead messaging channel; and
- maintain the protocol of eoc message exchange (re-send messages, abandon certain tasks, etc.).

The VME-O shall update and store the set of near-end test parameters (the ones that can be updated during Showtime) within 10 seconds after receiving the request to do so from the NMS.

The VME-O shall update and store the set of far-end test parameters (the ones that can be updated during Showtime) within 30 seconds after receiving the request to do so from the NMS.

11.2.2 eoc transmission protocol

A VTU invokes eoc communication with the VTU at the other end of the link by sending an eoc command message. The responding VTU, acting as a slave, shall acknowledge a command it has received correctly by sending a response, unless one is not required for the particular command type.

Furthermore, it shall perform the requested management function. Both VTUs shall be capable of sending eoc commands and responding to received eoc commands. The same eoc protocol format shall be used in both transmission directions. To send commands and responses over the line, the VME originates eoc messages. Each eoc message is a command, a command segment, a response, or a response segment. The VME sends each eoc message to the MPS-TC.

The MPS-TC encapsulates all incoming messages into HDLC format, as specified in clause 8.2.3. The length of any eoc message shall be less than or equal to 1024 octets, as described in clause 11.2.3.1.

Each command and the corresponding response are associated with a priority level specified in clause 11.2.3.1. To maintain priorities of eoc commands when sent over the link, the VME shall send messages to the MPS-TC via the γ_m interface in accordance with the priority levels of the commands (responses) carried by these messages, as specified in Table 11-1.

Table 11-1 – eoc message priority levels

Priority level	Associated time-out value	eoc command (response)
High	400 ms	Table 11-2, UTC (see clause 11.2.3.2)
Near High	For further study	Table 11-3
Normal	800 ms	Table 11-4
Low	1 s	Table 11-5

The VME shall send the eoc command only once and wait for a response, if one is required. No more than one command of each priority level shall be awaiting a response at any time. Upon reception of the response, a new command of the same priority level may be sent. If the command is segmented, all the segments of the command shall be sent and responses received before the next command is sent.

Accordingly, the VME shall send the message carrying a command or a segment of a command only once and wait for a response message. Upon reception of the response message, a new message may be sent. If a response to a particular message is not received within a specified time period (see Table 11-1), or is received incorrectly, a time-out occurs. After a time-out, the VME shall re-send the message up until REINIT_TIME_THRESHOLD seconds from the first time-out after which it shall abandon the message.

From all of the messages available for sending at any time, the VME shall always send the message with highest priority first. If a message with a higher priority than the one that is currently being sent becomes available for sending, the VME may abort sending the lower priority message (by setting the *Tx_Stop* signal, as specified in clause 8.2.4.1). The VME shall re-send the aborted message as the priority rule allows (i.e., when its priority level is the highest among all messages available for sending).

Messages of different priority have different time-out durations, as shown in Table 11-1, except for messages for which a response is not required and hence no timeout period is applicable. Time-outs shall be calculated from the instant the MPS-TC sends the last octet of the message until the instant the VME receives the first octet of the response message. Accordingly, the time-out timer shall be started by the *Sent* signal. If the VME detects an *Rx_RF* signal and a corresponding *Rx_PrF* signal within the relevant time-out value specified in Table 11-1, it shall set a time stamp for the preliminary arrival time of the expected response message, and then wait for the *Rx_Enbl* signal; otherwise the VME shall time-out for the expected response.

If the VME detects the *Rx_Enbl* signal in ≤300 ms after *Rx_RF* and *Rx_PrF* signals are set, the response message is considered to be received; otherwise, the VME shall consider the received *Rx_RF* and *Rx_PrF* signals as false, and shall delete the time stamp and wait for the next *Rx_RF* and *Rx_PrF* signals within the rest of the time-out value specified in Table 11-1.

The receiver uses the assigned value specified in clause 11.2.3.2 to determine the type and priority of the received eoc command (response).

11.2.3 eoc commands and responses

11.2.3.1 General

The first octet of a command (response) specifies the type of command (response). The second octet specifies the name of the command (response) for the specified type. Other octets carry the management data associated with the command (response).

The data values to be sent shall be mapped such that the LSB of data is mapped to the LSB of the corresponding octet of the command (response). Data values containing more than one octet shall be mapped with higher order octets preceding lower order octets. A vector of data values shall be mapped in order of the index, from the lowest index value to the highest.

If a specific command (response) is longer than 1024 octets, the VME shall segment it as specified in clause 11.2.3.3 so that the length of the eoc messages sent is shorter than *P* octets. The maximum length *P* of the message shall be based on the assigned message overhead data rate in the relevant transmission direction using the following equation:

$$P \leq \min\,(1024,\, 33 \times msg_p) \text{ octets}$$

where:

msg_p = message overhead data rate for latency path p in kbit/s (specified in clause 9.5.4).

NOTE – With the defined value of *P*, the transmission time of any eoc message will not exceed 270 ms (including 3% loss due to HDLC overhead and stuffing). This ensures that in all regular cases the VME will not be forced to stop sending a low-priority message in order to comply with the time-out requirements presented in Table 11-1. The VME should avoid long commands and responses.

11.2.3.2 Command and response types

With the exception of control parameter read, which is for further study, the VTU shall support all mandatory eoc command and response types specified in Table 11-2 (high priority commands), Table 11-3 (near high priority commands), Table 11-4 (normal priority commands) and Table 11-5 (low priority commands), and their associated commands and responses specified in clauses 11.2.3.3 to 11.2.3.11, inclusive. The VTU should reply with Unable-To-Comply (UTC) response on the optional commands that the VTU cannot recognize the assigned value for the command type. The UTC response shall include two octets: the first octet of the UTC shall be the same as the first octet of the received command, and the second octet shall be FF_{16}. The UTC is a high priority response.

NOTE – If the UTC response is not supported, the command will time out. This would reduce the efficiency of the eoc.

Table 11-2 – High priority commands and responses

Command type and assigned value	Direction of command	Command content	Response content	Support
On-line reconfiguration (OLR) 0000 0001_2	From the receiver of either VTU to the transmitter of the other	All the necessary PMD and PMS-TC control parameter values for the new configuration	Includes either a line signal marking the instant of re-configuration (Syncflag), or an OLR intermediate acknowledge (for segmented command), or an OLR command to defer or reject the proposed reconfiguration	See Table 11-6
Power Management 0000 0111_2	From either VTU to the other	High priority LPM commands defined in Annex E of [ITU-T G.998.4]	An acknowledgement or reject of high priority LPM command as defined in Annex E of [ITU-T G.998.4]	Optional

Table 11-3 – Near high priority commands and responses

Command type and assigned value	Direction of command	Command content	Response content	Support
Frequency synchronization 0101 0000_2	From VTU-O to VTU-R	The ToD phase difference value to run frequency synchronization: the ns_counter value of the RTC-O mod 125000 ns divided by 2, which shall be represented by a 16-bit value.	No response needed	Optional

Table 11-4 – Normal priority commands and responses

Command type and assigned value	Direction of command	Command content	Response content	Support
Diagnostic 0100 0001_2	From VTU-O to VTU-R	Request to run the self-test, or to update test parameters, or to start and stop transmission of corrupt CRC, or to start and stop reception of corrupt CRC	Acknowledgment	Mandatory
	From VTU-R to VTU-O	Request to update test parameters	Acknowledgment	Mandatory

Table 11-4 – Normal priority commands and responses

Command type and assigned value	Direction of command	Command content	Response content	Support
Time $0100\ 0010_2$	From VTU-O to VTU-R	Set or read out the time	Acknowledgment of the set time command, or a response including the time value	Mandatory
Inventory $0100\ 0011_2$	From either VTU to the other	Identification request, auxiliary inventory information request, and self-test results request	Includes the VTU equipment ID auxiliary inventory information, and self-test results	Mandatory
Management counter read $0000\ 0101_2$	From either VTU to the other	Request to read the counters	Includes all counter values	Mandatory
Clear eoc $0000\ 1000_2$	From either VTU to the other	Clear eoc command as defined in [ITU-T G.997.1]	Acknowledgment	Mandatory
Power Management $0000\ 0111_2$	From either VTU to the other	Proposed new power state, or normal priority LPM commands defined in Annex E of [TU-T G.998.4]	An acknowledgement to either reject or grant the new power state, or acknowledgement or reject of normal priority LPM command as defined in Annex E of [ITU-T G.998.4]	Mandatory
Non-standard facility (NSF) $0011\ 1111_2$	From either VTU to the other	Non-standard identification field followed by vendor proprietary content	An acknowledgment or a negative acknowledgment indicating that the non-standard identification field is not recognized	Mandatory
Control parameter read $0000\ 0100_2$	From either VTU to the other	For further study	For further study	Mandatory
Time synchronization $0101\ 0001_2$	From VTU-O to VTU-R	Includes the time stamps obtained by VTU-O to run time synchronization	Includes either the corresponding time stamp values of events t_2 and t_3 to accept the time synchronization (ACK) or a reject of the time synchronization command with a reason code	Optional
SAVN-Update $1010\ 0001_2$	From VTU-O to VTU-R	Request to update the showtime adaptive virtual noise (SAVN)	Acknowledgment	Optional

Table 11-5 – Low priority commands and responses

Command type and assigned value	Direction of command	Command content	Response content	Support
PMD test parameter read $1000\ 0001_2$	From either VTU to the other	The identification of test parameters for single read, or for multiple read, or for block read	Includes the requested test parameter values or a negative acknowledgment	See Tables 11-26 and 11-27
Impulse noise monitoring (INM) facility $1000\ 1001_2$	From VTU-O to VTU-R	Set or readout the INM data	An acknowledgment of the INM facility set command, or a response including the INM data	Optional
Non-standard facility (NSF) low priority $1011\ 1111_2$	From either VTU to the other	Non-standard identification field followed by vendor proprietary content	An acknowledgment or a negative acknowledgment indicating that the non-standard identification field is not recognized	Mandatory
Data gathering $1100\ 1100_2$	From VTU-O to VTU-R	Configure data gathering or request a number of data gathering records	ACK/NACK the data gathering configuration or transfer data gathering records	Optional

11.2.3.3 On-line reconfiguration (OLR) commands and responses

The VTU shall be capable of sending and receiving the OLR commands and responses listed in Tables 11-6 and 11-7, respectively, for the supported type(s) of OLR (see clause 13.1). Any OLR command specified in Table 11-6 may be initiated by either VTU. The responding VTU may either reject the initiator's request using responses listed in Table 11-7 with reason codes listed in Table 11-8, or positively acknowledge the initiator's request by transmitting a time marker for the reconfiguration. The time marker shall be communicated by transmission of a Syncflag (see clause 10.5.3). Changes may be requested concurrently by both VTUs; each transaction shall follow the procedure described in this clause.

The first octet of all OLR commands and responses shall be the assigned value for the OLR command type, as shown in Table 11-2. The remaining octets shall be as shown in Table 11-6 (for commands) and in Tables 11-7 and 11-8 (for responses). The octets of the OLR commands and responses shall be sent over the link as described in clause 11.2.3.1.

The list of parameters for any command in Table 11-6 shall be selected such that the length of the eoc message in octets (prior to HDLC encapsulation) does not exceed the maximum length P specified in clause 11.2.3.1. If more parameters are to be re-configured simultaneously, the initiator shall segment the Request command to meet the maximum message size. The number of segments shall not exceed 64. The multi-segment transmission is supported by the segment code (SC) octet in the Request command and by the intermediate acknowledge (IACK) octet in the response. The responding VTU shall send an IACK response after every intermediate segment has been received. After all segments have been received, the responding VTU shall send the Defer or Reject response with a reason code if the request cannot be processed, or send the time marker (Syncflag, see clause 10.5.3) to implement the request. The requesting VTU shall not send the next segment until it receives the IACK for the current segment. If an IACK for an intermediate segment is not received before the time-out, the requesting VTU may either re-send it or abandon the request. The responding VTU shall consider the OLR command abandoned if no more valid segments are received within 1 second of the last segment.

The two MSBs of the SC shall be set to 00_2 for intermediate segments and set to 11_2 for the last segment. The 6 LSBs shall contain the serial number of the segment starting from 000000_2. The SC octet of an IACK shall be the same as the SC octet of the acknowledged segment.

Table 11-6 – OLR commands sent by the initiating VTU

Name	Length (octets)	Octet number	Content	Support
Request Type 1	$5 + 4 \times N_f$ ($N_f \leq 128$)	2	04_{16} (Note 1)	Mandatory
		3 to 4	2 octets for the number of subcarriers N_f to be modified	
		5 to 4 + 4 × N_f	4 × N_f octets describing the subcarrier parameter field for each subcarrier	
		5 + 4 × N_f	1 octet for SC	
Request Type 2	For further study	2	05_{16} (Note 1)	For further study
		All others	Reserved by ITU-T	
Request Type 3 (SRA) (Note 6)	$5 + 7\ N_{LP} + 4\ N_f$ ($N_f \leq 128$)	2	06_{16} (Note 1)	Optional
		3 to 2 + 2 N_{LP}	2 × N_{LP} octets containing the new L_p values for each of the active latency paths (N_{LP} = number of active latency paths) (Notes 2 and 3)	
		3 + 2 N_{LP} to 2 + 4 N_{LP}	2 × N_{LP} octets containing the new D_p values for each of the active latency paths (N_{LP} = number of active latency paths) (Note 4)	
		3 + 4 N_{LP} to 2 + 5 N_{LP}	N_{LP} octets containing the new T_p values for each of the active latency paths (N_{LP} = number of active latency paths) (Notes 2, 3, 5)	
		3 + 5 N_{LP} to 2 + 6 N_{LP}	N_{LP} octets containing the new G_p values for each of the active latency paths (N_{LP} = number of active latency paths) (Notes 2, 3, 5)	
		3 + 6 N_{LP} to 2 + 7 N_{LP}	N_{LP} octets containing the new B_{p0} values for each of the active latency paths (N_{LP} = number of active latency paths) (Notes 2, 3, 5)	
		3 + 7 N_{LP} to 4 + 7 N_{LP}	2 octets for the number of subcarriers N_f to be modified	
		5 + 7 N_{LP} to 4 + 7 N_{LP} + 4 N_f	4 N_f octets describing the subcarrier parameter field for each subcarrier	
		5 + 7 N_{LP} + 4 N_f	1 octet for SC	

Table 11-6 – OLR commands sent by the initiating VTU

Name	Length (octets)	Octet number	Content		Support
Request Type 4 (SOS)	N_{TG}/2+11	2	07_{16} (Note 1)		Optional
		3	Message ID		
		4 to N_{TG}/2+3	$\Delta b(2)$	$\Delta b(1)$	
			$\Delta b(4)$	$\Delta b(3)$	
			…		
			$\Delta b(N_{TG})$	$\Delta b(N_{TG} - 1)$	
		N_{TG}/2+4 to N_{TG}/2+5	New value for L_0		
		N_{TG}/2+6 to N_{TG}/2+7	New value for L_1		
		N_{TG}/2+8 to N_{TG}/2+9	New value for D_0		
		N_{TG}/2+10 to N_{TG}/2+11	New value for D_1		
Request Type 5 (SRA/G.998.4)	See ITU-T G.998.4	2	08_{16} (Note 1)		Optional
		All others	Reserved for ITU-T G.998.4		
Request Type 6 (SOS/G.998.4)	See ITU-T G.998.4	2	09_{16} (Note 1)		Optional
		All others	Reserved for ITU-T G.998.4		

NOTE 1 – All other values for octet number 2 are reserved by ITU-T.

NOTE 2 – For this command, any change in L_p, T_p, G_p, and B_{p0} values shall be such that the length of the MDF (as defined in Table 9-8) remains unchanged for all active latency paths.

NOTE 3 – To keep the msg_p value within its valid range for relatively large changes of L_p, it may be necessary to change all of the T_p, G_p, and B_{p0} values.

NOTE 4 – If a change of D_p is not supported, the value of this parameter shall be identical to that currently used.

NOTE 5 – If a change of T_p, G_p and B_{p0} is not supported, the values of these parameters shall be identical to those currently used.

NOTE 6 – When N_{LP} = 2, the octets associated with latency path 0 are sent first.

The message ID identifies an SOS request. The message ID shall be an 8 bit wrap-around counter. The initial value of the message ID shall be set to 0 for the first SOS request after entering Showtime. When the SOS message is repeated, the same message ID shall be maintained for as long as the same request is sent. The next SOS request shall use a message ID that is incremented by 1.

The parameter N_{TG} is the number of SOS tone groups as specified for SOS in the O/R-PMS messages (see Table 12-56, Table 12-64).

$\Delta b(k)$ is the bit loading reduction in SOS tone group #k. These values shall be coded as 4-bit unsigned integers. The number of SOS tone groups shall be derived from the information exchanged in O-PMS and R-PMS. If that number is odd, the most significant four bits in byte #N_{TG}/2+3 shall be set to zero at the transmitter and ignored by the receiver.

An SOS request may be repeated before the SOS request time-out has expired. It is up to the VTU that receives the request to recognize that this is the same SOS request. Once its transmitter has acknowledged a request with given message ID (by sending a Syncflag), it shall ignore subsequent SOS requests with the same message ID.

For single latency with ROC mode, the L and D values for latency path #0 shall be ignored by the receiver.

Table 11-7 – OLR responses sent by the responding VTU

Name	Length (octets)	Octet number	Content	Support
Defer Type 1 request	3	2	81_{16} (Note)	Mandatory
		3	1 octet for reason code (Table 11-8)	
Reject Type 2 request	3	2	82_{16} (Note)	For further study
		3	1 octet for reason code (Table 11-8)	
Reject Type 3 request	3	2	83_{16} (Note)	Optional
		3	1 octet for reason code (Table 11-8)	
Reject Type 4 request	3	2	84_{16} (Note)	Optional
		3	1 octet for reason code (Table 11-8)	
Reject Type 5 request	3	2	85_{16} (Note)	Optional
		3	1 octet for reason code (Table 11-8)	
Reject Type 6 request	3	2	86_{16} (Note)	Optional
		3	1 octet for reason code (Table 11-8)	
IACK	3	2	$8B_{16}$ (Note)	Mandatory
		3	1 octet for SC	
NOTE – All other values for octet number 2 are reserved by ITU-T.				

Each subcarrier parameter field shall contain 4 octets formatted as [0000 iiii iiii iiii gggg gggg gggg bbbb] to convey the g_i (12 bits) and the b_i (4 bits) values of the subcarrier index *i* (12 bits). The subcarrier index *i* shall be coded in the four LSBs of the first octet and the entire second octet of the subcarrier field. The LSBs of the subcarrier index *i* shall be contained in the second octet. The g_i shall be contained in the third octet and the four MSBs of the fourth octet. The LSBs of g_i shall be contained in the fourth octet. The b_i shall be contained in the four LSBs of the fourth octet.

Table 11-8 – Reason codes for OLR responses

Reason	Octet value	Applicable to Defer Type 1	Applicable to Reject Type 2	Applicable to Reject Type 3	Applicable to Reject Type 4
Busy	01_{16}	yes	yes	yes	no
Invalid parameters	02_{16}	yes	yes	yes	yes

Upon sending an OLR command, the initiator shall await a response. The OLR response may be deferring or rejecting the reconfiguration, or it may be a Syncflag indicating when the reconfiguration shall take effect. If the initiator receives an OLR response to defer or reject the change, it shall abandon the last requested OLR command. A new command may be initiated immediately, including the command abandoned, rejected or deferred earlier.

NOTE 1 – In the case of reason code 02_{16}, repeating of the OLR request is not expected to be helpful.

NOTE 2 – When an OLR command has been sent, the initiator has no means to cancel the command. The initiator needs to wait for a response or a time-out before it can send a different OLR command. For example, if the SOS triggering conditions become active when there is a pending bitswap or SRA, the SOS request needs to be delayed until full execution or time-out of the bitswap/SRA procedure.

Upon reception of an OLR command, the responder shall send either an OLR response to defer or to reject the reconfiguration, or a Syncflag that indicates when the reconfiguration shall take effect. After sending the Syncflag, the responder shall reconfigure the affected PMD, PMS-TC, and TPS-TC functions starting from the tenth symbol in the next DMT superframe, as described in clause 13.3. The responder may defer or reject the OLR request; in this case it shall supply a reason code from those specified in Table 11-8.

Upon reception of the Syncflag, the initiator shall reconfigure the affected PMD or PMS-TC functions starting from the tenth DMT symbol in the next DMT superframe, as described in clause 13.3.

11.2.3.4 Diagnostic commands and responses

The Diagnostic commands shall be used to control the VTU diagnostic capabilities defined in this clause. The Diagnostic commands shown in Table 11-9 may be initiated only by the VTU-O. The Diagnostic commands shown in Table 11-10 may be initiated only by the VTU-R. The responses are shown in Table 11-11. All Diagnostic commands and responses shall consist of two or three octets. The first octet shall be the assigned value for the Diagnostic command type, as shown in Table 11-4. The second and subsequent octets shall be as shown in Tables 11-9 and 11-10 for commands and in Table 11-11 for responses. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-9 – Diagnostic commands sent by the VTU-O

Name	Length (Octets)	Octet number	Content
Perform self-test	2	2	01_{16} (Note)
Update Test Parameters	2	2	02_{16} (Note)
Start TX Corrupt CRC	2	2	03_{16} (Note)
End TX Corrupt CRC	2	2	04_{16} (Note)
Start RX Corrupt CRC	2	2	05_{16} (Note)
End RX Corrupt CRC	2	2	06_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-10 – Diagnostic commands sent by the VTU-R

Name	Length (Octets)	Octet number	Content
Update test Parameters	2	2	02_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-11 – Diagnostic responses sent by the VTU

Name	Length (Octets)	Octet number	Content
Self-test Acknowledge (VTU-R only)	3	2	01_{16} (Note)
		3	1 octet for the minimum time in seconds the VTU-O shall wait before requesting the self-test result
ACK (VTU-O and VTU-R)	2	2	80_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

A Diagnostic command may be sent at any time during showtime, including immediately following the end of the initialization procedure. In all cases, reception of a Diagnostic command shall be acknowledged to the initiator (by an ACK or by a self-test Acknowledge response).

NOTE – A negative acknowledge (NACK) is not used for Diagnostic commands.

11.2.3.4.1 Perform self-test

Upon reception of the Perform self-test command, the VTU-R shall respond with a self-test Acknowledge, which indicates the minimum amount of time that the VTU-O shall wait before requesting the results of the self-test. Further, the VTU-R shall perform the self-test and generate the self-test result. The self-test procedure is vendor discretionary, but it shall not interfere with the functions of the VTU-R, shall not impact the status of the connection, and its duration shall not exceed 255 s. The VTU-R shall obtain and store the result of the self-test within the number of seconds indicated in the self-test Acknowledge response. The indicated amount of time shall be an integer between 1 s and 255 s.

The self-test results may be accessed using the Inventory command defined in clause 11.2.3.6. The length of the self-test results shall be 4 octets. The first octet (including the MSB) shall be 00_{16} if the self-test passed and 01_{16} if it failed. The meaning of "failure" is vendor discretionary. The contents of the three other octets are vendor discretionary.

11.2.3.4.2 Update Test Parameters

Upon reception of the Update Test Parameters command, the requested VTU shall send the ACK response and update the test parameter set defined in clause 11.4.1. All test parameters that can be updated during showtime shall be updated and stored within 10 s after the request is received. Upon reception of the ACK response, the requesting VTU shall wait at least 10 s before sending the PMD Test Parameter Read commands defined in clause 11.2.3.11 to access the test parameter values defined in clause 11.4.1.

The test parameter values relating to the most recent initialization procedure shall no longer be accessible through the Test Parameter Read commands within 10 s after the Update Test Parameters command was received. They may be discarded by the responding VTU immediately upon reception of the Update Test Parameter command.

11.2.3.4.3 Start/End transmitter Corrupt CRC

Upon reception of the Start Transmitter (TX) Corrupt CRC command, the VTU-R shall send the ACK response and its PMS-TC shall generate a corrupted CRC value in all transmitted latency paths until cancelled by the End TX Corrupt CRC command. A corrupted CRC value is any one that does not correspond to the CRC procedure specified in clause 9.5.2.3. The Start TX Corrupt CRC command shall affect only the CRC value transmitted by the VTU-R; the PMS-TC function of the VTU-O shall not be affected by this command.

Upon reception of the End TX Corrupt CRC command, the VTU-R shall send the ACK response and its PMS-TC shall generate CRC values in all latency paths as specified in clause 9.5.2.3. The End TX Corrupt CRC command shall not affect the PMS-TC function of the VTU-R if the Start TX Corrupt CRC command has not been sent earlier.

11.2.3.4.4 Start/End receiver Corrupt CRC

Upon reception of the Start receiver (RX) Corrupt CRC command, the VTU-R shall send the ACK response. Upon reception of this ACK response by the VTU-O, its PMS-TC function shall generate a corrupted CRC value in all transmitted latency paths until cancelled by the End RX corrupt CRC command. A corrupted CRC value is any one that does not correspond to the CRC procedure specified in clause 9.5.2.3. The Start RX Corrupt CRC command shall affect only the CRC value transmitted by the VTU-O; the PMS-TC function of the VTU-R shall not be affected by this command.

Upon reception of the End RX Corrupt CRC command, the VTU-R shall send the ACK response. Upon reception of this ACK response, the PMS-TC function at the VTU-O shall generate CRC values in all latency paths as specified in clause 9.5.2.3. The End RX Corrupt CRC command shall not affect the PMS-TC function of the VTU-O if the Start RX Corrupt CRC command has not been sent earlier.

NOTE – The Start RX Corrupt CRC command may be used in conjunction with the Transmit Corrupt CRC command (either previously or subsequently) so that CRC values are set corrupted in both directions of transmission.

11.2.3.5 Time commands and responses

Both VTUs shall maintain timers to update performance monitoring counters as described in [ITU-T G.997.1]. The time commands shall be used to synchronize timers at both ends of the link. The timers shall have an accuracy of ±100 ppm or better.

NOTE – The counters defined in [ITU-T G.997.1] should be updated each time the time counter contains a time value that is an integer multiple of 15 minutes (e.g., 1:00:00, 3:15:00, 15:30:00, 23:45:00).

The time commands are shown in Table 11-12, and may only be initiated by the VTU-O. The VTU-R shall reply using one of the responses shown in Table 11-13. The first octet of all time commands and responses shall be the assigned value for the time command type, as shown in Table 11-4. The remaining octets shall be as specified in Tables 11-12 and 11-13 for commands and responses, respectively. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-12 – Time commands sent by the VTU-O

Name	Length (Octets)	Octet number	Content
Set Time	10	2	01_{16} (Note)
		3 to 10	8 octets for time value formatted as HH:MM:SS per [ISO 8601]
Read Time	2	2	02_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-13 – Time responses sent by the VTU-R

Name	Length (Octets)	Octet number	Content
ACK	2	2	80_{16} (Note)
Time	10	2	82_{16} (Note)
		3 to 10	8 octets for time value formatted as HH:MM:SS per [ISO 8601]
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Upon reception of the Set Time command, the VTU-R shall send the ACK response, and set its timer to the value contained in the message.

Upon reception of the Read Time command, the VTU-R shall send the Time response that includes the current value of the VTU-R timer.

11.2.3.6 Inventory commands and responses

The Inventory commands shall be used to determine the identification and capabilities of the VTU at the far end. The Inventory commands shown in Table 11-14 may be initiated by either VTU. The Inventory responses shall be as shown in Table 11-15. The first octet of all Inventory commands and responses shall be the assigned value for the Inventory command type, as shown in Table 11-4. The second octet of the Inventory commands shall be as specified in Table 11-14. The second octet (ACK) and all following octets of the Inventory responses shall be as specified in Table 11-15. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-14 – Inventory commands sent by the requesting VTU

Name	Length (Octets)	Octet number	Content
Identification request	2	2	01_{16} (Note)
Auxiliary Inventory Information request	2	2	02_{16} (Note)
Self-test results request	2	2	03_{16} (Note)
Initialization Flags request	2	2	04_{16} (Note)
Initialization Flags Reset request	2	2	05_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-15 – Inventory responses sent by the responding VTU

Name	Length (Octets)	Octet number	Contents
ACK (Identification)	58	2	81_{16} (Note)
		3 to 10	8 octets of vendor ID
		11 to 26	16 octets of version number
		27 to 58	32 octets of serial number
ACK (Auxiliary Inventory Information)	variable	2	82_{16} (Note)
		3 to 10	8 octets of vendor ID
		11 +	Multiple octets of auxiliary inventory information
Self-test Results	6	2	83_{16} (Note)
		3 to 6	4 octets of self-test results
Initialization Flags	3	2	84_{16} (Note)
		3	1 octet with the value of the initialization flags.
Initialization Flags Reset	3	2	85_{16} (Note)
		3	1 octet with the value of the initialization flags before the reset.
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Upon reception of one of the Inventory commands, the VTU shall send the corresponding response. Any function of either the requesting or the responding VTU shall not be affected by the command.

The vendor ID in the response identifies the system integrator and shall be formatted according to the vendor ID of [ITU-T G.994.1]. In the context of this request, the system integrator usually refers to the vendor of the smallest field-replaceable unit; thus, the vendor ID in the response may not be the same as the vendor ID indicated during the ITU-T G.994.1 handshake phase of initialization.

The VTU-O version number shall be as defined in clause 7.4.5 of [ITU-T G.997.1].

The VTU-R version number shall be as defined in clause 7.4.6 of [ITU-T G.997.1].

The VTU-O serial number shall be as defined in clause 7.4.7 of [ITU-T G.997.1].

The VTU-R serial number shall be as defined in clause 7.4.8 of [ITU-T G.997.1].

The auxiliary inventory information shall be assigned with respect to the same system integrator as contained in the vendor ID. The syntax of this field is beyond the scope of this Recommendation.

The self-test results response shall contain the results from the most recent self-test procedure, initiated either at power-up or by the eoc command Perform self-test. The results shall be formatted as defined in clause 11.2.3.4.1.

The eoc commands Initialization Flags request and the Initialization Flags Reset request shall only be supported from the VTU-O to the VTU-R. The responses to those commands are optional.

The Initialization Flags and the Initialization Flags Reset response shall contain the current value of the initialization flags. The following initialization flags are defined:

- The "previous-loss-of-power" (PLPR) flag: This flag shall be set to 1 after a power-up of the VTU-R due to an interruption in the VTU-R electrical supply (mains) power. The flag shall be set to 0 after sending the Initialization Flags Reset response;
- The "previous host re-init" (PHRI) flag: This flag shall be set to 1 after a power-up of the VTU-R triggered by the CPE host. The flag shall be set to 0 after sending the Initialization Flags Reset response.

The value of the initialization flags shall be formatted as 1 octet [0000 00ba] where "a" is the value of the PLPR flag and "b" is the value of the PHRI flag.

11.2.3.7 Management counter read commands and responses

The Management counter read request command shall be used to retrieve the current value of certain management counters maintained by the far-end VTU in accordance with [ITU-T G.997.1]. The Management counter read request command is shown in Table 11-16, and may be initiated by either VTU and is used to request the values of the counters. The response shall be as shown in Table 11-17. The first octet of the command and response shall be the assigned value for the Management counter read command type, as shown in Table 11-4. The second octet of the command shall be as shown in Table 11-16. The second and all following octets of the response shall be as shown in Table 11-17. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-16 – Management counter read commands sent by the requesting VTU

Name	Length (Octets)	Octet number	Content
Request	2	2	01_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-17 – Management counter read responses sent by the responding VTU

Name	Length (Octets)	Octet number	Content
ACK	variable	2	81_{16} (Note 1)
		3 to 2 + 4 × (2 × N_{LP} + 5)	Octets for all of the PMS-TC counter values (Note 2)
		3 + 4 × (2 × N_{LP} + 5) and above	Octets for all of the TPS-TC counter values (Note 2)
NOTE 1 – All other values for octet number 2 are reserved by ITU-T. NOTE 2 – N_{LP} is the number of enabled latency paths.			

Upon reception of the management counter read request command, the VTU shall send the response. Any function of either the requesting or the responding VTU shall not be affected by the command.

The management counter values shall be derived according to [ITU-T G.997.1] from locally generated defects and anomalies defined within clause 11.3. The parameters shall be transferred in the order (top to bottom) defined in Table 11-18. The TPS-TC anomaly definitions and relevant management counters are dependent upon the TPS-TC type and shall be as defined in Annex L. All counter values are defined as 32-bit counters and shall be mapped to the response in order of most significant to least significant octet. No octets shall be inserted into the response for latency paths and TPS-TC functions that are currently disabled.

The counters shall be reset at power-up, and shall not be reset upon a link state transition, and shall not be reset upon read. The time periods when the VTU is powered but not in the showtime state shall be counted as unavailable seconds (see clause 7.2.1.1.5 of [ITU-T G.997.1]).

Table 11-18 – VTU management counters

PMS-TC counters
Counter of the FEC-0 anomalies
Counter of the FEC-1 anomalies
Counter of the CRC-0 anomalies
Counter of the CRC-1 anomalies
FEC errored seconds counter
Errored seconds counter
Severely errored seconds counter
los errored seconds counter
Unavailable errored seconds counter
TPS-TC counters
Counters for TPS-TC #0
Counters for TPS-TC #1

NOTE – The VTU-O should respond to the request from the NMS to read the values of management counters. It is left to the implementations to store and update the counters as necessary for accurate error monitoring and reporting.

11.2.3.8 Clear eoc commands and responses

The Clear eoc Request command may be used by the ITU-T G.997.1 function to transfer management octets between the EIA and the VTU-R and from one VTU to another (see clause 6 of [ITU-T G.997.1]). The Clear eoc Request command is shown in Table 11-19 and may be initiated by either VTU. The responses shall be as shown in Table 11-20. The first octet of either the command or a response shall be the assigned value for the Clear eoc command type shown in Table 11-4. The subsequent octets of the command shall be as shown in Table 11-19. The subsequent octets of the responses shall be as shown in Table 11-20. The octets shall be sent using the format described in clause 11.2.3.1.

NOTE – In accordance with [ITU-T G.997.1], the information payload of the Clear eoc message does not exceed 510 octets. Therefore, the length of either a Clear eoc Request command or a response does not exceed 516 octets.

Table 11-19 – Clear eoc commands sent by the initiating VTU

Name	Length (Octets)	Octet number	Content
Request	variable	2	01_{16} (Note)
		3 +	the information payload of the clear eoc message to be delivered to the far end
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-20 – Clear eoc responses sent by the responding VTU

Name	Length (Octets)	Octet number	Content
ACK	2	2	80_{16} (Note)
NACK	3	2	81_{16} (Note)
		3	04_{16} (Note)
NOTE – All other values for octet numbers 2 and 3 are reserved by ITU-T.			

Upon reception of the Clear eoc Request command, the VTU shall respond with an acknowledgement (ACK) and deliver the received clear eoc message to the local ITU-T G.997.1 management function transparently, with the original formatting used by the ITU-T G.997.1 management function of the initiating VTU. The VTU may instead respond with a negative acknowledge (NACK) including the Not Supported (value 04_{16}) reason code, indicating that the received clear eoc message cannot be delivered to the ITU-T G.997.1 management function (because the ITU-T G.997.1 management function may not support clear eoc messages; see clause 6 of [ITU-T G.997.1]). Other reason codes are for further study.

11.2.3.9 Power management commands and responses

The Power Management L3 Request command shall be used to propose a power management transition to link state L3. The Power Management L3 Request command is shown in Table 11-21 and may be initiated by either VTU. The responses shall be as shown in Table 11-22. The first octet of either the command or a response shall be the assigned value for the Power Management command type, as shown in Table 11-4. The remaining octets shall be as shown in Tables 11-21 and 11-22 for commands and responses, respectively.

Table 11-21 – Power management commands sent by the initiating VTU

Name	Length (Octets)	Octet number	Content
L3 Request	3	2	01_{16} (Note)
		3	03_{16} (Note)
NOTE – All other values for octet numbers 2 and 3 are reserved by ITU-T.			

Table 11-22 – Power management responses sent by the responding VTU

Name	Length (Octets)	Octet number	Content
Grant	2	2	80_{16} (Note)
Reject	3	2	81_{16} (Note)
		3	1 octet for reason code
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Reason codes associated with the power management commands are shown in Table 11-23.

Table 11-23 – Reason codes for power management commands

Reason	Octet value
Busy	01_{16}
Invalid	02_{16}
State not desired	03_{16}

11.2.3.9.1 L3 Request by VTU-R

Upon receipt of the L3 Request command, the responding VTU-O shall send either the Grant or Reject response. The proposed link state shall be formatted as 03_{16} for the L3 link state. If any other link state is received, the Reject response shall be sent with the reason code 02_{16}.

The VTU-O may reject a request to move to link state L3 using reason code 01_{16} because it is temporarily too busy, or reject it using code 03_{16} because it has local knowledge that the L3 state is not desired at this time. Upon receipt of the L3 Request command, the VTU-O may immediately start the protocol to request a transition to the L3 state.

If the VTU-R receives the Grant response, the VTU-R shall stop transmitting. When the VTU-O observes the stopped transmission, it shall also stop transmitting.

11.2.3.9.2 L3 Request by VTU-O

Upon receipt of the L3 Request command, the responding VTU-R shall send either the Grant or Reject response. The proposed link state shall be formatted as 03_{16} for the L3 link state. If any other link state is received, the Reject response shall be sent with the reason code 02_{16}.

The VTU-R may reject a request to move to link state L3 using reason code 01_{16} because it is temporarily too busy, or reject it using code 03_{16} because it has local knowledge that the L3 state is not desired at this time. Upon receipt of the L3 Request command, the VTU-R may immediately start the protocol to request a transition to the L3 state.

If the VTU-O receives the Grant response, the VTU-O shall stop transmitting. When the VTU-R observes the stopped transmission, it shall also stop transmitting.

11.2.3.10 Non-standard facility commands and responses

The non-standard facility (NSF) commands may be used to exchange vendor-discretionary information between the VTUs. The NSF Request command is shown in Table 11-24 and may be initiated by either VTU to request the non-standard information. The responses shall be as shown in Table 11-25. The first octet of either the command or a response shall be the assigned value for the NSF command type, as shown in Table 11-4 for normal priority NSF commands, or in Table 11-5 for low priority NSF commands. The remaining octets of normal priority and low priority commands shall be as shown in Table 11-24. The second octet of normal priority and low priority responses shall be as shown in Table 11-25. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-24 – NSF commands sent by the requesting VTU

Name	Length (Octets)	Octet number	Content
Request	variable	2	01_{16} (Note)
		3 to 8	6 octets of NSF identifier field
		9 +	multiple octets of NSF message field
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table 11-25 – NSF responses sent by the responding VTU

Name	Length (Octets)	Octet number	Content
ACK	2	2	80_{16} (Note)
NACK	2	2	81_{16} (Note)
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Upon reception of the NSF Request command, the VTU shall respond with an acknowledgement (ACK) to indicate that both the NSF identifier field and the message field are recognized, or respond with a negative acknowledgement (NACK) if either the NSF identifier field or NSF message field is not recognized.

The combination of the NSF identifier field and NSF message field corresponds to a non-standard information block as defined in Figure 11 of [ITU-T G.994.1] (without the length-indicator octet). The NSF identifier field shall consist of 6 octets. The first 2 octets shall be a country code, and the remaining 4 octets shall be a provider code as specified by the country. Both values shall be set as defined in [ITU-T T.35]. The NSF message field contains vendor-specific information. The syntax of the NSF message field shall be as defined in Figure 11 of [ITU-T G.994.1] (without the length-indicator octet).

11.2.3.11 PMD Test Parameter Read commands and responses

The PMD Test Parameter Read commands shall be used to retrieve the values of the PMD test parameters that are specified in clause 11.4.1 and maintained by the far-end VTU. The PMD Test Parameter Read commands are shown in Table 11-26, and may be initiated by either VTU. The responses shall be as shown in Table 11-27. The first octet of all PMD Test Parameter Read commands and responses shall be the assigned value for the PMD Test Parameter Read command type, as shown in Table 11-5. The subsequent octets of the commands shall be as shown in Table 11-26. The subsequent octets of the responses shall be as shown in Table 11-27. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-26 – PMD test parameter read commands sent by the requesting VTU

Name	Length (octets)	Octet number	Content	Support
Single Read	2	2	01_{16} (Note 1)	Mandatory
Next Multiple Read	2	2	03_{16} (Note 1)	Mandatory
	4	2	04_{16} (Note 1)	Mandatory

Table 11-26 – PMD test parameter read commands sent by the requesting VTU

Name	Length (octets)	Octet number	Content	Support
Multiple Read		3 to 4	2 octets describing the subcarrier group index	
Block Read	6	2	05_{16} (Note 1)	Mandatory
		3 to 4	2 octets describing the start subcarrier group index	
		5 to 6	2 octets describing the stop subcarrier group index	
Vector Block Read	7	2	06_{16} (Note 1)	Optional
		3	1 octet describing the type of test parameter to read (Note 2) 01_{16}: Channel transfer function Hlog(*f*) per subcarrier group 03_{16}: Quiet Line Noise PSD QLN(*f*) per subcarrier group 04_{16}: Signal to noise ratio SNR(*f*) per subcarrier group.	
		4 to 5	2 octets describing the start subcarrier group index	
		6 to 7	2 octets describing the stop subcarrier group index	
Scalar Read	3	2	07_{16} (Note 1)	Optional
		3	1 octet describing the type of scalar test parameters to be read (Note 2) 21_{16} to 28_{16}: the parameter index to read according to the ID of Table 11-28.	
NOTE 1 – All other values for octet number 2 are reserved by the ITU-T. NOTE 2 – All other values for octet number 3 are reserved by the ITU-T.				

Table 11-27 – PMD test parameter read responses sent by the responding VTU

Name	Length (octets)	Octet number	Content	Support
Single Read ACK	42 (Note 1)	2	81_{16} (Note 2)	Mandatory
		3 to 42	Octets for the test parameters arranged for the single read format	
Multiple Read ACK	12 (Note 1)	2	82_{16} (Note 2)	Mandatory
		3 to 12	Octets for the test parameters arranged for the multiple read format	
NACK	2	2	80_{16} (Note 2)	Mandatory
Block Read ACK	Parameter-dependent (Note 1)	2	84_{16} (Note 2)	Mandatory
		3 +	Octets for the test parameters arranged for the block read format	
Vector Block Read ACK	Parameter-dependent (Note 1)	2	86_{16} (Note 2)	Optional
		3 +	Octets for the test parameters arranged for the block read format	

Table 11-27 – PMD test parameter read responses sent by the responding VTU

Name	Length (octets)	Octet number	Content	Support
Scalar Read ACK	Parameter-dependent (Note 1)	2	87_{16} (Note 2)	Optional
		3 +	Octets for the test parameters arranged for the scalar read format	
NOTE 1 – Message length equals 2 octets plus the length shown in Table 11-28. NOTE 2 – All other values for octet number 2 are reserved by the ITU-T.				

Table 11-28 – PMD test parameter ID values and length of responses

Test parameter ID (Note 1)	Test parameter name	Length for Single Read (octets)	Length for Multiple Read (octets)	Length for Block Read or Vector Block Read (octets)	Length for Scalar Read (octets)	Support
01_{16}	Channel transfer function Hlog(f) per subcarrier group	N/A	4	2 + (stop subcarrier group index – start subcarrier group index + 1) × 2 (Note 2)	N/A	Mandatory
03_{16}	Quiet line noise PSD QLN(f) per subcarrier group	N/A	3	2 + (stop subcarrier group index − start subcarrier group index + 1) (Note 2)	N/A	Mandatory
04_{16}	Signal-to-noise ratio SNR(f) per subcarrier group	N/A	3	2 + (stop subcarrier group index − start subcarrier group index + 1) (Note 2)	N/A	Mandatory
21_{16}	Loop attenuation (LATN)	2 × 5	N/A	N/A	2 × 5	Mandatory
22_{16}	Signal attenuation (SATN)	2 × 5	N/A	N/A	2 × 5	Mandatory
23_{16}	Signal-to-noise ratio margin (SNRM) and SNRM-pb	2 × 6	N/A	N/A	2 × 6	Mandatory
24_{16}	Attainable net data rate ATTNDR (basic method)	4	N/A	N/A	4	Mandatory
24_{16}	Attainable net data rate ATTNDR (improved method)	8	N/A	N/A	8	Optional

Table 11-28 – PMD test parameter ID values and length of responses

Test parameter ID (Note 1)	Test parameter name	Length for Single Read (octets)	Length for Multiple Read (octets)	Length for Block Read or Vector Block Read (octets)	Length for Scalar Read (octets)	Support
25_{16}	Near-end actual aggregate transmit power ACTATP	2	N/A	N/A	2	Mandatory
26_{16}	Far-end actual aggregate transmit power ACTATP	2	N/A	N/A	2	Mandatory
27_{16}	Far-end actual impulse noise protection INP_act	N/A	N/A	N/A	2	Optional
28_{16}	Far-end actual signal-to-noise ratio margin for the robust overhead channel SNRM-ROC	N/A	N/A	N/A	2	Optional
NOTE 1 – All other Test parameter ID values are reserved by the ITU-T. NOTE 2 – Since the number of subcarriers, *G*, in the subcarrier group (see clause 11.4.1) may be different for QLN, Hlog, and SNR, the values of QLN, Hlog and SNR communicated by Multiple Read, Block Read, or Vector Block Read for the same subcarrier group index may correspond to different subcarrier indices. The subcarrier index for each parameter equals *G* × subcarrier group_index, where the value of *G* is as defined in Table 11-42 of clause 11.4.1 (for showtime) and subcarrier group index = 0 to 511.						

Upon reception of a PMD Test Parameter Read command, the responding VTU shall send the corresponding response. If the format of the Test Parameter Read command is incorrect, the VTU shall respond with the negative acknowledge (NACK). Any function of either the requesting or the responding VTU shall not be affected.

The Single Read command shall be used to retrieve all test parameters with ID values from 21_{16} to 26_{16} inclusive. In response to a Single Read command, the values for the test parameters (one value per parameter) shall be transferred in numerically increasing order of the parameter ID shown in Table 11-28. The format of the octets for each parameter shall be as specified in clause 11.4.1. Values formatted as multiple octets shall be mapped to the response in order of most significant to least significant octet. The loop attenuation (LATN), signal attenuation (SATN) and signal-to-noise ratio margin (SNRM) format shall include five 2-octet values intended for 5 potentially available frequency bands for each transmission direction. The 2-octet values shall be sent in the order shown in Table 11-29. The value 00_{16} shall be used to indicate the disabled bands. Octets indicated as reserved shall be set to ZERO in the transmitter and ignored by the receiver. The SNRM test parameter shall, in addition to all SNRM-pb values (clause 11.4.1.1.6.3), include the overall SNRM value (clause 11.4.1.1.6.2). The first 2-octet value is the overall SNRM, followed by the five 2-octet values of the SNRM-pb as specified in Table 11-29. For the ATTNDR, the use of either the basic or the improved method is configured during initialization (see clause 11.4.1.1.7). The ATTNDR test parameter is specified in Table 11-30.

Table 11-29 – Order for sending LATN, SATN and SNRM-pb parameters

Octet number	Upstream direction	Downstream direction
1	US0	DS1
2		
3	US1	DS2
4		
5	US2	DS3
6		
7	US3	DS4
8		
9	US4	Reserved
10		

Table 11-30 – ATTNDR test parameter

Octet number	Basic method	Improved method
1-4	*ATTNDR*	*ATTNDR*
5	N/A	Reserved and set to 00_{16}
6	N/A	$ATTNDR\_INP\_act_0$
7	N/A	Reserved and set to 00_{16}
8	N/A	$ATTNDR\_delay\_act_0$
NOTE – The format of the fields is defined in clause 11.4.1.1.7.		

A Scalar Read command shall be used to retrieve a single test parameter. Support of this read command is optional. The ID of the test parameter to retrieve shall be indicated in the third octet of the read command as specified in Table 11-26. In response to a Scalar Read command, the VTU shall send the value of the test parameter if this command and the test parameter are supported by the VTU; otherwise the VTU shall send a NACK. The format of the octets for each parameter value shall be as described in clause 11.4.1. Values formatted as multiple octets shall be mapped to the response in order of most significant to least significant octet. The format of the LATN, SATN, SNRM and ATTNDR shall be identical to the format used in Single Read Command. The Far-end actual impulse noise protection (ID=27_{16}) shall include two 1-octet values and be sent in the order shown in Table 11-31. The value FF_{16} shall be used to indicate the disabled bearers.

Table 11-31 – Order for sending far-end actual impulse noise protection parameters

Octet number	Parameter
1	INP_act for bearer channel 0
2	INP_act for bearer channel 1

Multiple Read and Next Multiple Read commands shall be used to retrieve test parameters of one subcarrier group. In response to a Multiple Read or Next Multiple Read command, the VTU shall send information for test parameters with ID 01_{16}, 03_{16}, and 04_{16} associated with the indicated subcarrier group. The Multiple Read command contains the index of the requested subcarrier group (see Table 11-26). If a Next Multiple Read command is to be sent, it shall only be sent after a Multiple Read command. In response to each subsequent Next Multiple Read command, the subcarrier group index shall be incremented by one. If the subcarrier group index exceeds 511 (see clause 11.4.1), the response shall be a NACK. The values of the PMD parameters per subcarrier group shall be inserted into the message in numerical order of the parameter ID shown in Table 11-28. The format of the octets for each parameter shall be as described in clause 11.4.1. Values that are formatted as multiple octets shall be mapped to the response in order of most significant to least significant octet.

A Block Read command shall be used to retrieve test parameters over a range of subcarrier groups. In response to a Block Read command, the VTU shall send information for test parameters with ID 01_{16}, 03_{16}, and 04_{16} associated with the specified block of subcarrier groups. For test parameters specified per subcarrier group, all values for subcarrier groups with indices from #start to #stop are transferred in a single response. If the subcarrier group index exceeds 511, the response shall be a NACK. The values of the PMD parameters per subcarrier group shall be inserted into the message in increasing order of the parameter ID shown in Table 11-28. The format of the octets for each parameter value shall be as described in clause 11.4.1. Values formatted as multiple octets shall be mapped to the response in order of most significant to least significant octet. The number of octets in a Block Read command shall not exceed the maximum length *P* of the eoc message specified in clause 11.2.3.1.

A Vector Block Read command shall be used to retrieve a single test parameter over a range of subcarrier groups. Support of this read command is optional. The ID of the test parameter to retrieve shall be indicated in the third octet of the read command as specified in Table 11-26. In response to a Vector Block Read command, the VTU shall send information for the test parameter associated with the specified block of subcarrier groups if this command is supported by the VTU; otherwise the VTU shall send a NACK. All values for subcarrier groups with indices from #start to #stop are transferred in a single response. If the subcarrier group index exceeds 511, the response shall be a NACK. The format of the octets for each parameter value shall be as described in clause 11.4.1. Values formatted as multiple octets shall be mapped to the response in order of most significant to least significant octet.

When transferring values of the channel transfer function Hlog(*f*), the quiet line noise QLN(*f*), and the signal-to-noise ratio SNR*(f)*, the measurement time shall be included in the response for each test parameter. The ACK (see Table 11-27) is followed by the HLOG*(f)* measurement time and the value *m* (see clause 11.4.1.1.1), followed by the QLN*(f)* measurement time and value *n* (see clause 11.4.1.1.2), followed by the SNR*(f)* measurement time and value SNR (see clause 11.4.1.1.3), respectively. The measurement time shall be included only once in a response to a Block Read or Vector Block Read command and shall be included for each test parameter in each response to a Multiple Read or Next Multiple Read command.

The values of some test parameters are represented using fewer bits than contained in the corresponding field defined for the response in Table 11-28. In the case that the field has more than one octet, the bits shall be mapped to the LSBs of the multi-octet field in the response. Unused MSBs in the multi-octet field shall be set to ZERO for unsigned quantities and to the value of the sign bit for signed quantities.

11.2.3.12 Control Parameter Read commands and responses

Control parameter read commands are for further study.

11.2.3.13 INM facility commands and responses

A VTU that supports the impulse noise monitoring (INM) facility shall maintain INM counters to measure the impulse noise, as described in [ITU-T G.997.1]. The INM facility commands shall be used to update and read the INM parameters at the VTU-R.

The INM facility command shall be used also to retrieve the current value of the INM counters maintained by the far-end VTU in accordance with [ITU-T G.997.1].

The INM facility commands are described in Table 11-32, and may only be initiated by the VTU-O. The VTU-R shall reply using one of the responses shown in Table 11-33. The first octet of all INM facility commands and responses shall be the assigned value for the INM facility command type, as shown in Table 11-5. The remaining octets shall be as specified in Table 11-32 and Table 11-33 for commands and responses, respectively. The octets shall be sent using the format described in clause 11.2.3.1.

Table 11-32 – INM facility commands sent by the VTU-O

Name	Length (Octets)	Octet number	Content
Read INM counters	2	2	02_{16}
Set INM parameters	6 or 7	2	03_{16}
		3 to 6 or 7	4 or 5 octets of INM parameters: see Table 11-36
Read INM parameters	2	2	04_{16}
All other values for octet number 2 are reserved by ITU-T.			

Table 11-33 – INM facility responses sent by the VTU-R

Name	Length (Octets)	Octet number	Content
ACK	3	2	80_{16}
		3	1 octet INM acceptance code: see Table 11-34
NACK	2	2	81_{16}
INM counters	107	2	82_{16}
		3 to 2 + 4 × (17+1+8)	Octets for all of the INM counter values: see Table 11-35
		107	1 octet INMDF
INM parameter s	6 or 7	2	84_{16}
		3 to 6 or 7	4 or 5 octets of INM parameters: see Table 11-36
All other values for octet number 2 are reserved by ITU-T.			

Upon reception of any INM facility command, the VTU-R shall send NACK in response if it does not support the INM procedure or the INM command is invalid. Upon reception of an INM facility set INM parameters command, the VTU-R shall send the ACK in response if it does support the INM procedure.

In case all INM parameter values listed in the set INM parameters command are valid and supported by the VTU-R, the VTU-R shall accept all of the INM parameters contained in the command. The INM acceptance code (see Table 11-34) will indicate that the parameters are accepted. If, for any of the INM parameters, the value in the command is different from the value in active use by the INM, the VTU-R shall activate the new INM parameter values and reset the counters less than 1 second after sending the ACK.

In case any INM parameter values listed in the set INM parameters command is valid but not supported by the VTU-R, the VTU-R shall not accept any of the INM parameters and shall not reset the counters.

Upon reception of the INM facility read INM parameters command, the VTU-R shall send the INM parameters response that includes the current value of the VTU-R INM parameters.

Table 11-34 – VTU-R INM acceptance code

Name	Octet No.	Content
ACC-INM_INPEQ_MODE	3	80_{16}: value for INM_INPEQ_MODE accepted
NACC-INM_INPEQ_MODE	3	81_{16}: value for INM_INPEQ_MODE not supported

Upon reception of the INM facility read INM counters command, the VTU shall send the INM counters response, which includes the INMDF. Any function of either the requesting or the responding VTU shall not be affected by this command.

The INM counter values shall be derived according to [ITU-T G.997.1] from locally generated defects and anomalies defined in clause 11.3. The parameters shall be transferred in the order (top to bottom) defined in Table 11-35. All counter values are defined as 32-bit counters and shall be mapped to the response in order of most significant to least significant octet.

The INM counters shall be reset at power-up, and shall not be reset upon any link state transition, and shall not be reset upon read. They shall be reset at activation of the new INM parameter values. The reset value is zero. The INM counters and the procedure to update the counters shall work continuously and independently of other (proprietary or non-proprietary) features, e.g., the INM shall work in the presence of OLR and without interruption.

NOTE – The VTU-O should respond to the request from the NMS to read the values of INM counters. It is left to the implementations to store and update the counters as necessary for accurate monitoring and reporting.

Table 11-35 – VTU-R INM counters

INM counters
Counter of the $INMAINPEQ_1$ anomalies
Counter of the $INMAINPEQ_2$ anomalies
...
Counter of the $INMAINPEQ_{16}$ anomalies
Counter of the $INMAINPEQ_{17}$ anomalies
Counter of the $INMAIAT_0$ anomalies
Counter of the $INMAIAT_1$ anomalies
...
Counter of the $INMAIAT_6$ anomalies
Counter of the $INMAIAT_7$ anomalies
Counter of the INMAME anomalies

The VTU-R shall set the INM default flag (INMDF) to ONE whenever all active INM parameters are equal to the default values. The VTU-R shall set the INM default flag (INMDF) to ZERO whenever any active INM parameter is different from the default value.

The INM parameter values shall be transferred in the order defined in Table 11-36 and mapped in order of most significant to least significant octet.

Table 11-36 – VTU-R INM parameters

Octet No.	INM parameter
3-4	2 octets: • The 9 LSBs are INMIATO • The 4 MSBs are INMIATS
5	1 octet: INMCC
6	1 octet: INM_INPEQ_MODE
7	1 octet: INM_INPEQ_FORMAT This octet is present only if supported by both VTU-O and VTU-R (see Tables 12-49 and 12-60).

11.2.3.14 Frequency Synchronization command

The VTU-O shall be capable of sending frequency synchronization command through the eoc if this mechanism is selected during initialization. The command is only in one direction, from VTU-O to VTU-R, and there is no response required.

The ToD phase difference value and corresponding t_1 event number shall be encapsulated in an *eoc* message as follows:

Table 11-37 – Frequency Synchronization command sent by the VTU-O

Name	Length (Octets)	Octet number	Content
ToD phase difference	5	2	03_{16} (Note)
		3	1 octet representing the index of the t_1 event
		4 to 5	2 octets representing the ToD phase difference in units of 2 nanoseconds
NOTE – All other values for octet number 2 are reserved by ITU-T.			

If the VTU-R selects during initialization that the ToD phase difference values shall be transported through the eoc, then the VTU-O shall generate a ToD phase difference eoc message into the nearly-high priority queue for each super-frame. For each super-frame, the VTU-O shall send the most recent phase difference eoc message in the queue and shall discard older phase difference eoc messages. If the VTU-R selects during initialization to transport the ToD phase difference values through the OH frame, then the VTU-O shall not send phase difference eoc messages.

11.2.3.15 Time Synchronization command and responses

The VTU shall be capable of sending and receiving the Time Synchronization commands and responses listed in Table 11-38 (command sent by VTU-O) and Table 11-39 (response sent by VTU-R), respectively. The Timestamp command specified in Table 11-38 shall only be sent by the VTU-O. The Timestamp response specified in Table 11-39 shall only be sent by the VTU-R. The VTU-R may reject to run time synchronization procedure using responses listed in Table 11-39 with reason codes listed in Table 11-40, or positively acknowledge by transmitting an ACK response.

The first octet of all Time Synchronization commands and responses shall be the assigned value for the Time Synchronization command type, as shown in Table 11-4. The remaining octets shall be as shown in Tables 11-38 and 11-39. The octets of the Time Synchronization commands and responses shall be set over the link as described in clause 11.2.3.1.

Table 11-38 – Time Synchronization commands sent by the VTU-O

Name	Length (Octets)	Octet number	Content
ToD(t_1) ToD(t_4) Timestamps	26	2	01_{16} (Note 1)
		3 to 4	2 octet for the index of time stamp ToD(t_1) in units of a super-frame.
		5 to 10	6 octets describing the integer portion of the timestamp ToD(t_1) in units of seconds.
		11 to 14	4 octets describing the fractional portion of the timestamp ToD(t_1) in units of nanoseconds. (Note 2)
		15 to 16	2 octets for the index of time stamp ToD(t_4) in units of a superframe.
		17 to 22	6 octets describing the integer portion of the timestamp ToD(t_4) in units of seconds.
		23 to 26	4 octets describing the fractional portion of the timestamp ToD(t_4) in units of nanoseconds. (Note 2)
NOTE 1 – All other values for octet number 2 are reserved by ITU-T. NOTE 2 – The nanosecond portion is always less than 10^9.			

The octets for the index of time stamp ToD(t_1) contain the value of the downstream super-frame counter when ToD(t_1) is taken by the VTU-O (i.e., at the t_1 event). This value shall be a multiple of 16. The octets for the index of time stamp ToD(t_4) contain the value of the upstream super-frame counter when ToD(t_4) is taken by the VTU-O (i.e., at t_4 event). The difference between the downstream super-frame counter at t_1 event and the upstream superframe counter at t_4 event shall be constant over Showtime (i.e., the pairing of the t_1 and t_4 events shall not change over Showtime). The t_1 event and the t_4 event shall be less than 1 superframe apart (i.e., 64.25 ms on the PMD sampling clock timebase if the CE length corresponds to m = 5, see clause 10.4.4).The ToD(t_1), ToD(t_2), ToD(t_3), and ToD(t_4) time stamps are described as two parts. One is the integer portion of the timestamp in units of seconds and the other is the fractional portion of the timestamp in units of nanoseconds. The ToD(t_1), and ToD(t_4) time stamps shall represent the time offset between the current time of the Real-time Clock RTC-O at the VTU-O (i.e., the time elapsed since the epoch) at the t_1 and t_4 events respectively. The ToD(t_2), and ToD(t_3) timestamps shall represent the time of the Real-time Clock RTC-R at the VTU-R, (i.e., the time elapsed since the epoch) at the t_2 *and* t_3 events respectively. The epoch shall be the same for the Real-time Clock RTC-O and the Real-time Clock RTC-R, where this common epoch is set over the γ-O reference point.

NOTE – If at the t_1 event the Real-time Clock RTC-O shows +2.000000001 seconds have elapsed since the epoch, this is represented in the ToD(t_1) timestamp by seconds = 0x0000 0000 0002 and nanoseconds = 0x0000 0001. The epoch may be locally set by the DSLAM or may be an absolute instant in time. For example, if the epoch is the PTP epoch, this means that time-of-day = 1 January 1970 00:00:02.000000001.

Table 11-39 – Time Synchronization responses sent by the VTU-R

Name	Length (Octets)	Octet number	Content
ToD(t_2) ToD(t_3) Timestamps (ACK)	26	2	81_{16} (Note 1)
		3 to 4	2 octets for the index of t_2 time stamp.
		5 to 10	6 octets describing the integer portion of the timestamp in units of seconds.
		11 to 14	4 octets describing the fractional portion of the timestamp in units of nanoseconds. (Note 2)
		15 to 16	2 octets for the index of t_3 time stamp.
		17 to 22	6 octets describing the integer portion of the timestamp in units of seconds.
		23 to 26	4 octets describing the fractional portion of the timestamp in units of nanoseconds. (Note 2)
Reject	3	2	82_{16} (Note 1)
		1	1 octet for reason code (see Table 11-40)
NOTE 1 – All other values for octet number 2 are reserved by ITU-T. NOTE 2 – The nanosecond portion is always less than 10^9.			

Table 11-40 – Reason codes for Time Synchronization responses

Reason	Octet value
Busy	01_{16}
Invalid parameters	02_{16}
t_2 and t_3 timestamps no longer available at the VTU-R	03_{16}
Still acquiring ToD frequency synchronization	04_{16}

The Timestamp command is used to send time stamps ToD(t_1) and ToD(t_4) from the VTU-O to the VTU-R. Upon reception of a Timestamp command, the VTU-R shall either send the time stamps ToD(t_2) and ToD(t_3) in an ACK response to indicate that the time synchronization procedure will be performed with the ToD(t_1), ToD(t_2), ToD(t_3) and ToD(t_4) timestamps, or send a reject response with a reason code from those specified in Table 11-40.

The VTU-R shall store the t_2 (and related t_3) values for at least the three most recent downstream (and related upstream) reference samples with a t_2 event count that is a multiple of 16. The VTU-O should send the time synchronization command soon enough after the t_1 and t_4 events to assure the related t_2 and t_3 timestamps are still available at the VTU-R.

If the VTU-R accepts the Timestamp command, the Timestamp response is used to send timestamps ToD(t_2), and ToD(t_3) from the VTU-R to the VTU-O. The ToD(t_1), ToD(t_2), ToD(t_3) and ToD(t_4) timestamps (in conjunction with other information) may be used at the network side to e.g., compensate for propagation delay asymmetry. At the customer premises side, propagation delay asymmetry shall not be compensated for. Other uses of the response reported time stamp values at the network side are for further study.

11.2.3.16 Data gathering commands and responses

This clause defines the commands and responses to support the data gathering function defined in clause 11.5. If the VTU-R supports the data gathering function, then the VTU-R shall keep a buffer of data gathering records, with each data gathering record relating to one of the event types listed in Table 11-43. Data gathering records shall be transmitted from the VTU-R to the VTU-O in eoc messages upon request of the VTU-O. The VTU-O sends a data gathering command (to configure data gathering or to request a number of data gathering records, see Table 11-40.1). The VTU-R shall respond with a data gathering response (to ACK/NACK the data gathering configuration or to transfer data gathering records, see Table 11-40.2). The VTU-O shall acknowledge (ACK) the response and may simultaneously request more data gathering records. The VTU-R buffer shall be transferred to the VTU-O through a sequence of command and response messages until all records in the VTU-R buffer, up to *act_logging_depth_reporting_R* records, are reported.

Table 11-40.1 – Data gathering commands sent by the VTU-O

Name	Length (Octets)	Octet number	Content
Configure data gathering		2 3-4 5 6 7 … 7+*Ntyp*	00_{16} *act_logging_depth_reporting_R* *logging_reporting_newer_first* number of event types configured (*Ntyp*) *logging_ depth_event_percentage_R_1* … *logging_ depth_event_percentage_R_Ntyp*
Request data gathering records	1 1 1 1	2 3 4 5	01_{16} ACK_ID Sequence_ID Number of event records to report in response message (*Nreq*)

Table 11-40.2 – Data gathering response sent by the VTU-R

Name	Length (Octets)	Octet number	Content
Configure data gathering ACK	4	4 5-6	80_{16} *logging_depth_R*
Configure data gathering NACK	4	4 5-6	81_{16} *logging_depth_R*
Response with data gathering records	10 (if *Nrep*=0) up to 778 (if *Nrep*=128)	2 3 4-7 8-9 10 11 ... 10+6×*Nrep*	90_{16} Sequence_ID Current VTU-R timestamp Number of event records in the buffer not yet ACKed (*Nnack*) Number of event records reported in this message (*Nrep*) 6×*Nrep* octets with the data gathering records

11.2.3.16.1 Data gathering configuration

The control parameter *act_logging_depth_reporting_R* is defined in clause 11.5. The control parameter *logging_report_newer_first* indicates the order in which the VTU-R shall report the records in the VTU-R buffer. If set to TRUE, then the VTU-R shall report the record related to the most recent event first. If set to FALSE, then the VTU-R shall report the record related to the least recent event first. The control parameter *logging_report_newer_first* shall have a value equal to the CO-MIB configuration parameter LOGGING_REPORT_NEWER_FIRST (see clause 7.3.6.5 of [ITU-T G.997.1]).

The control parameters $logging\_depth\_event\_percentage\_R_i$ are defined in clause 11.5. The $logging\_depth\_event\_percentage\_R_i$ is configured for event type ID = 1 up to event type ID = *Ntyp*, with event type IDs as listed in Table 11-43. For all event types with ID > *Ntyp*, the VTU-R shall set $logging\_depth\_event\_percentage\_R_i = 0$.

If the VTU-R supports the last received data gathering configuration, then the VTU-R shall send a configure data gathering ACK response. If the VTU-R does not support the last received data gathering configuration (e.g., because *act_logging_depth_reporting_R* is higher than *logging_depth_R*), then the VTU-R shall send a configure data gathering NACK response. Upon sending an ACK, the VTU-R shall apply the last received data gathering configuration. Upon receiving a NACK, the VTU-O shall send a new configure data gathering command.

If a new data gathering configuration is applied with an *act_logging_depth_reporting_R* larger than or equal to the actual number of records in the VTU-R buffer, then existing records shall stay in the VTU-R buffer. Otherwise, the VTU-R may apply a vendor discretionary method (including reset of the entire VTU-R buffer) to discard records in the buffer.

If a new data gathering configuration is applied with an *act_logging_depth_reporting_O* larger than or equal to the actual number of records in the VTU-O buffer, then existing records shall stay in the VTU-O buffer. Otherwise, the VTU-O may apply a vendor discretionary method (including reset of the entire VTU-O buffer) to discard records in the buffer.

11.2.3.16.2 Transfer of the VTU-R buffer from the VTU-R to the VTU-O

The VTU-R buffer is transferred to the VTU-O through a sequence of command and response messages until all records in the VTU-R buffer, up to *act_logging_depth_reporting_R* records, are reported.

The VTU-O allocates a sequence ID to each request data gathering records command message. The first command message in a sequence of command and response messages shall have Sequence_ID = 00_{16}. If the response message to a command message is received by the VTU-O, then the next command message shall have the Sequence_ID incremented by one. If the response message to a command message is not received, then the next command message shall be a retransmission of the previous command message or a command message with sequence_ID = 00_{16}. If a sequence of more than 128 command and response messages is needed to transfer the whole VTU-R buffer, then the sequence_ID following sequence_ID = $7F_{16}$ shall be sequence_ID = 01_{16}.

In the response with data gathering records message, the VTU-R shall copy the same sequence_ID allocated by the VTU-O to the request data gathering records command message.

In the command message with Sequence_ID = 00_{16}, the VTU-O shall set the ACK_ID to 00_{16}.

When receiving a command message with Sequence_ID = 00_{16}, the VTU-R shall ignore the ACK_ID in this command message. Before sending a response message, the VTU-R shall mark all records in the VTU-R buffer (except dummy records) as "not yet ACKed". In the response message (with Sequence_ID = 00_{16}), the value *Nnack* indicate the number of records in the VTU-R buffer. With a VTU-R buffer full condition, *Nnack* = *act_logging_depth_reporting_R*.

When receiving a command message with Sequence_ID > 00_{16}, and before sending a response message, the VTU-R shall mark the set of records which were reported in the response message with sequence_ID equal to *ACK_ID* as having been ACKed. The VTU-R shall keep the ACKed records in the VTU-R buffer (i.e., not delete the ACKed records from the VTU-R buffer).

NOTE 1 – The records are not marked as ACKed until the ACK_ID is received because eoc messages (requests or responses) can get lost.

The command message shall include the number of records requested to be reported (*Nreq*). The value of *Nreq* shall be at most 128. If the VTU-O chooses to use multiple eoc command/response messages to transfer the whole VTU-R buffer, then the value of *Nreq* may be less than *act_logging_depth_reporting_R*.

The response message shall include the number of records stored in the VTU-R buffer that are not yet ACKed by the VTU-O (*Nnack*), so the VTU-O knows how many more records are to be transferred.

The response message shall include the number of records reported in the response message (*Nrep*). The value *Nrep* shall be the minimum of *Nreq* and *Nnack*. If *Nrep* < *Nreq*, then *Nrep* = *Nnack*. If all records in the VTU-R buffer have been ACKed, then *Nrep* = *Nnack* = 0.

If *logging_report_newer_first* is FALSE, then the response message shall include the *Nrep* least recent records in the VTU-R buffer that have not yet been ACKed by the VTU-O. The response message shall include the *Nrep* records in ascending order of timestamp, with the least recent record starting at octet 11.

If *logging_report_newer_first* is TRUE, then the response message with sequence_ID = 0 shall include the *Nrep* most recent records in the VTU-R buffer that have not yet been ACKed by the VTU-O. The response messages with sequence_ID > 0 shall include the *Nrep* most recent records in the VTU-R buffer that have not yet been ACKed by the VTU-O and have a timestamp older than the timestamps of the records ACKed by the VTU-O in the command message. The response message shall include the *Nrep* records in descending order of timestamp, with most recent record starting at octet 11.

The response message includes the current VTU-R timestamp. This allows the VTU-O to offset all the VTU-R timestamps with a fixed offset so that they are stored in the CO-MIB in the same NTP format, with nearly the same timing, as the VTU-O timestamps. This also limits the time error due to VTU-R clock drift to the accumulated error between the time the event is recorded and the time that it is transferred to the VTU-O.

Upon receiving the response message containing the last of *act_logging_depth_reporting_R* records to be transferred, or upon receiving a response message with *Nnack* = *Nrep*, the VTU-O shall send the last command message in the sequence of command and response messages that is an acknowledgement of the last transferred *Nrep* records, and has *Nreq* = 0, denoting no further records are to be transferred. Only in this last command message in a sequence of command and response messages shall *Nreq* = 0. All preceding command messages shall have *Nreq* > 0.

NOTE 2 – If no new records are generated at the VTU-R during the eoc transfer, then the entire VTU-R buffer is transferred to the VTU-O.

NOTE 3 – If new event records are generated at the VTU-R during the eoc transfer, leading to already transferred records being removed from the VTU-R buffer due to buffer management rules, and *logging_report_newer_first* is FALSE, then the VTU-R buffer will not be entirely transferred to the VTU-O because *act_logging_depth_reporting_R* records will have been ACKed by the VTU-O while *Nnack* is still larger than zero at the VTU-R.

NOTE 4 – If new event records are generated at the VTU-R during the eoc transfer, leading to not yet transferred records being removed from the VTU-R buffer due to buffer management rules, and *logging_report_newer_first* is TRUE, then the VTU-R buffer will not be entirely transferred to the VTU-O because *Nnack* will be zero at the VTU-R before *act_logging_depth_reporting_R* records have been ACKed by the VTU-O.

11.2.3.16.3 Write of the transferred VTU-R buffer to the CO-MIB

If *logging_report_newer_first* is TRUE, then (see Notes 3, 4 and 5 below):

- When the VTU-O sends a data gathering command message with sequence_ID = 00_{16}, the VTU-O shall reset the EVENT_TRACE_BUFFER_R to all zero values by setting all records to be dummy records, where a dummy record is defined as a record which has all fields set to the value 0;
- The VTU-O shall write the block of records contained in the data gathering response message with Sequence_ID = 00_{16} in the EVENT_TRACE_BUFFER_R in the locations corresponding to the most recent block of records;
- The VTU-O shall write blocks of records contained in subsequent data gathering response messages in the EVENT_TRACE_BUFFER_R in the locations corresponding to the next most recent block of records;
- The VTU-O may choose to update the EVENT_TRACE_BUFFER_R with every received data gathering response message, or may choose to update after having received several data gathering response messages;

  NOTE 1 – Subsequent readings of the EVENT_TRACE_BUFFER_R over the Q-interface while a transfer of the VTU-R buffer is ongoing may show a steadily increasing number of records in the CO-MIB element.
- The order of the records in the EVENT_TRACE_BUFFER_R shall correspond to the order of the records inside a received data gathering response message (i.e., no reordering inside a block), and concatenated in the order of subsequently received data gathering response messages (i.e., no reordering of blocks).

If *logging_report_newer_first* is FALSE, then (see Notes 4 and 5 below):

- Then the VTU-O sends a data gathering command with sequence_ID = 00_{16}, then the EVENT_TRACE_BUFFER_R is not reset;
- The VTU-O shall write the block of records contained in the data gathering response message with Sequence_ID = 00_{16} in the EVENT_TRACE_BUFFER_R in the locations corresponding to the least recent block of records and thereby overwrite the existing records at those locations;
- The VTU-O shall write blocks of records contained in subsequent data gathering response messages in the EVENT_TRACE_BUFFER_R in the locations corresponding to the next least recent block of records and thereby overwrite the existing records at those locations;
- The VTU-O may choose to update the EVENT_TRACE_BUFFER_R with every received data gathering response message, or may choose to update after having received several data gathering response messages;

  NOTE 2 – Subsequent readings of the EVENT_TRACE_BUFFER_R over the Q-interface while a transfer of the VTU-R buffer is ongoing may show some duplicate records.
- The order of the records in the EVENT_TRACE_BUFFER_R shall correspond to the order of the records inside a received data gathering response message (i.e., no reordering inside a block), and concatenated in the order of subsequently received data gathering response messages (i.e., no reordering of blocks).

  NOTE 3 – Setting *logging_report_newer_first* to TRUE allows the rapid transfer of a relatively small number of recent event records with a relatively small number of data gathering commands.

  NOTE 4 – If no new event records are generated at the VTU-R during the transfer of the VTU-R buffer, then the entire VTU-R buffer is transferred into EVENT_TRACE_BUFFER_R.

NOTE 5 – If new records are generated at the VTU-R during the transfer of the VTU-R buffer, then the CO-MIB parameter EVENT_TRACE_BUFFER_R may not contain some records that are present in the VTU-R buffer at the end of the transfer, may not contain some records that were present in the VTU-R buffer at the start of the transfer, and may contain records that are not present in the VTU-R buffer at the end of the transfer because of buffer management.

11.2.3.17 SAVN-Update command and responses

The VTU-O shall send the SAVN-Update command to communicate the updated value of the downstream transmitter referred showtime adaptive virtual noise (TXREFSAVNds) to the VTU-R. Upon reception of the SAVN-Update command, the VTU-R shall send a SAVN-Update response.

The first octet of both the command and the response indicates the SAVN-Update command type, as defined in Table 11-4 (normal priority). The other octets of the SAVN-Update command are defined in Table 11-40.3. The SAVN-Update response shall be a simple acknowledgement that indicates the serial number of the update, as defined in Table 11-40.4. The SAVN-Update command shall be initiated by the VTU-O only. The SAVN-Update response shall be sent by the VTU-R only.

Table 11-40.3 – SAVN-Update command (sent by the VTU-O)

Name	Length (Octets)	Octet number	Content
SAVN-Update	7+(stop sub-carrier group index − start sub-carrier group index + 1)	2	09_{16} (Note 1)
		3	One octet containing the sequence number of the update (Note 2)
		4 to 5	2 octets describing the start subcarrier group index
		6 to 7	2 octets describing the stop subcarrier group index
		8 to 7+(stop subcarrier group index − start subcarrier group index + 1)	The transmitter-referred SAVN PSD, TXREFSAVNds
NOTE 1 – All other values for octet numbers 2 are reserved by ITU-T. NOTE 2 – The sequence number of the update shall be "1" for the first SAVN-Update command, and shall be incremented by 1 for each subsequent SAVN-Update command, wrapping around to "0" after "255".			

TXREFSAVNds indicates the PSD of the SAVN that shall be taken into account when determining bit loading and active tone set using SNRM_MODE = 5 (see clause 11.4.1.1.6.1.5). In case SAVN is not used, both the start subcarrier group index and the stop subcarrier group index shall be equal to 0 and the TXREFSAVNds field shall be of zero length.

TXREFSAVNds($k \times G \times \Delta f$) shall be the value of the transmitter referred SAVN PSD for subcarriers with indices from $k \times G$ to $((k+1) \times G) - 1$, inclusive. It shall be represented as an 8 bit unsigned integer $savn(k)$, where $k = 0$ to 511. The value of TXREFSAVNds($k \times G \times \Delta f$) shall be defined as TXREFSAVNds($k \times G \times \Delta f$) $= -23 - (savn(k)/2)$ dBm/Hz. This data format supports a TXREFSAVNds(f) granularity of 0.5 dB with a valid range of values for TXREFSAVNds(f) from −150 (coded 254) to −23 (coded 0) dBm/Hz. The value coded 255 is reserved for future use by ITU-T.

The group size for TXREFSAVNds shall be $G = 8$.

Table 11-40.4 – SAVN-Update response (sent by the VTU-R)

Name	Length (Octets)	Octet number	Content
SAVN-ACK	4	2	$0A_{16}$ (Note 1)
		3	One octet containing the sequence number of the acknowledged SAVN-Update command
		4	One octet indicating whether or not an SRA will follow (Note 1): 00_{16} SRA will follow 01_{16} No SRA will follow
SAVN-NACK	3	2	$0B_{16}$ (Note 1)
		3	Reason code (Note 1): 00_{16} Invalid parameter set 01_{16} Violation of initialization policy
NOTE 1 – All other values for this octet are reserved by ITU-T.			

The VTU-R shall compute the bit loading and framing parameters corresponding to the updated value of the TXREFSAVNds. The updated value of TXREFSAVNds may require a modification of the bit loading, or framing parameters, or both, with the aim of maintaining operation of the line within the boundaries set at initialization (e.g., SNRM boundaries). If no modification of the bit loading, or framing parameters is required, then the VTU-R shall reply with the SAVN-ACK response, indicating no SRA will follow. If modification of the bit loading or framing parameters is required and the required modification of the bit loading or framing parameters is within the boundaries set at initialization (e.g., related to ETR), then the VTU-R shall reply with the SAVN-ACK response indicating an SRA will follow and shall initiate an SRA (Request Type 3 for ITU-T G.993.2 and Request Type 5 for ITU-T G.998.4) within 1 second after sending the SAVN-ACK response. If the required modification cannot be made within the boundaries set at initialization (e.g., related to ETR), then the VTU-R shall reply with the SAVN-NACK response using a reason code 01_{16}, which means that the requested SAVN PSD cannot be applied within the initialization policy.

Upon receiving a SAVN-NACK response, the VTU-O may either send another SAVN PSD update or re-initialize the line with new boundaries for performance parameters, which can fit the required SAVN settings.

VTU-O shall not send SAVN-Update command when the VTU-R is involved in an OLR procedure associated with change of the bit loading or framing parameters in the downstream direction.

11.2.4 Indicator bits (IB)

The IB are used to send the far-end anomalies and defects specified in Table 11-41. Sending IB is mandatory, both upstream and downstream. The IB shall be set to ZERO if in the active state. Mapping of the IB to the overhead channel shall be as specified in clause 9.5.2.2.

Table 11-41 – Content of IB

IB	Description	Reference
los	Loss of signal defect	See clause 11.3.1.3
rdi	Remote defect indication defect	See clause 11.3.1.4
lpr	Loss of power primitive	See clause 11.3.3.1
TIB#0-1 to TIB#0-4	Four indicator bits reserved for the TPS-TC serving bearer #0	See Annex L
TIB#1-1 to TIB#1-4	Four indicator bits reserved for the TPS-TC serving bearer #1	See Annex L

11.3 OAM primitives

Among the standard OAM primitives, this Recommendation specifies only anomalies and defects. The system shall use the corresponding failure specifications of [ITU-T G.997.1].

Both the near-end and the far-end primitives shall be represented at the VTU-O; representation of the far-end anomalies and defects at the VTU-R is optional.

OAM primitives shall not be generated for the overhead-only latency path (ROC).

11.3.1 Line-related primitives

Line-related primitives represent anomalies, defects, and other primitives related to PMD and PMS-TC sublayers.

11.3.1.1 Near-end anomalies

- Forward error correction (*fec-p*): This anomaly occurs when a received FEC codeword in the latency path #*p* indicates that errors have been corrected. This anomaly is not asserted if errors are detected and are not correctable.
- Cyclic redundancy check (*crc-p*): This anomaly occurs when a received CRC byte for the latency path #*p* is not identical to the corresponding locally generated CRC byte.
- Rate adaptation upshift (*rau*)*:* For further study.
- Rate adaptation downshift (*rad*): For further study.
- Loss-of-power interruption (*lpr_intrpt*): Excluding re-initializations triggered by the VTU-O host, this anomaly occurs upon entry into showtime when at least one of the following conditions is met:
  – a far-end loss-of-power (*flpr*) primitive was declared before the exit from showtime; or
  – the PLPR flag (retrieved from the VTU-R) is set to ONE at the entry into showtime (see clause 11.2.3.6).

  This anomaly is only defined at the VTU-O;
- Host-Reinit interruption (*hri_intrpt*)*:* Excluding re-initializations triggered by the VTU-O host, this anomaly occurs upon entry into showtime when the PHRI flag (retrieved from the VTU-R) is set to ONE at the entry into Showtime (see clause 11.2.3.6).

  This anomaly is only defined at the VTU-O.
- Spontaneous interruption (*spont_intrpt*)*:* Excluding re-initializations triggered by the VTU-O host, this anomaly occurs upon entry into showtime after a successful full initialization, if the following conditions are met:
  – the time between the exit from showtime of the VTU-O and the first successful reception of a [ITU-T G.994.1] message is less than 23 seconds; and

– neither an *lpr_intrpt* anomaly nor an *hri_intrpt* anomaly occurs at the entry into showtime.

The PLPR flag and the PHRI flag are optional at the VTU-R. If not supported by the VTU-R, then the VTU-O shall consider these flags as set to ZERO.

This anomaly is only defined at the VTU-O.

11.3.1.2 Far-end anomalies

- Far-end forward error correction (*ffec-p*): This anomaly occurs when an *fec-p* anomaly detected at the far end is reported. This anomaly terminates when the received report on the *fec-p* anomaly is terminated.
- Far-end block error (*febe-p*): This anomaly occurs when a *crc-p* anomaly detected at the far end is reported. This anomaly terminates when the received report on the *crc-p* anomaly is terminated.

11.3.1.3 Near-end defects

- Loss of signal (*los*): A reference power is established by averaging the VDSL2 receive power over a 0.1 s period and over a subset of subcarriers used for Showtime, and a threshold shall be set 6 dB below this level. An *los* occurs when the level of the VDSL2 receive power averaged over a 0.1 s period and over the same subset of subcarriers is lower than the threshold, and terminates when this level, measured in the same way, is at or above the threshold. The subset of subcarriers is implementation dependent.
- Severely errored frame (*sef*): This defect occurs when the content of two consecutively received sync symbols does not correlate with the expected content over a subset of the subcarriers. An *sef* terminates when the content of two consecutively received sync symbols correlates with the expected content over the same subset of the subcarriers. The correlation method, the selected subset of subcarriers, and the threshold for declaring these defect conditions are vendor discretionary.
- Loss of margin (*lom*)*:* This defect occurs when the signal-to-noise ratio margin (SNRM), see clause 11.4.1.1.6, observed by the near-end receiver is below the minimum signal-to-noise ratio margin (MINSNRM, see clause 12.3.5.2.1.1) and an increase of SNRM is no longer possible within the far-end aggregate transmit power and transmit PSD level constraints. This defect terminates when the SNRM is above the MINSNRM. The SNRM measurement update rate shall be at least once every 5 seconds.

11.3.1.4 Far-end defects

- Far-end loss of signal (*los-fe*): This defect occurs when an *los* detected at the far end is reported in at least 4 of 6 consecutively received far-end *los* indicator reports. An *los-fe* terminates when fewer than two far-end *los* indicators are reported out of 6 consecutively received reports.
- Remote defect indication (*rdi*): This defect occurs when an *sef* detected at the far end is reported. An *rdi* terminates when the received report on *sef* is terminated, i.e., when the value of the corresponding IB is reset to ONE.
- Far-end loss of margin (*lom-fe*): This defect occurs when the SNRM, see clause 11.4.1.1.6, at the far-end receiver, retrieved by the near-end transmitter is below the minimum signal-to-noise ratio margin (MINSNRM, see clause 12.3.5.2.1.1) and an increase of SNRM is no longer possible within the near-end aggregate transmit power and transmit PSD level constraints. This defect terminates when the SNRM is above the MINSNRM.

11.3.1.5 Near-end OLR primitives

Successful SOS (*success_SOS*): This primitive occurs each time a successful SOS procedure is performed through the near-end initiating an OLR type 4 or 6 for SOS (see Table 11-6 and clause 12.1.4).

Successful autonomous SRA (*success_SRA*): This primitive occurs each time the bit loading on the data symbols is changed through the SRA initiated by the near-end (using Request Type 3, see Table 11-6).

11.3.1.6 Far-end OLR primitives

Successful SOS (*success_SOS-fe*): This primitive occurs each time a successful SOS procedure is performed through the far-end initiating an OLR type 4 or 6 for SOS (see Table 11-6 and clause 12.1.4).

Successful autonomous SRA (*success_SRA_FE*): This primitive occurs each time the bit loading on the data symbols is changed through the SRA initiated by the far-end (using Request Type 3, see Table 11-6).

11.3.2 Path-related primitives

Path-related primitives are defined separately for each path, terminated by the corresponding TPS-TC. The primitives for each TPS-TC (ATM, PTM, STM, etc.) shall be represented by relevant OAM indicators specified for this protocol.

11.3.2.1 Anomalies and defects for ATM transport

The specified set of anomalies and defects for the ATM transport shall be supported by the ATM-TC. In the case of multiple bearer channels, the corresponding ATM-TCs shall be represented by independent sets of indicators. The anomalies and defect indicators shall comply with clause L.2.

11.3.2.2 Anomalies and defects for STM transport

The specified set of anomalies and defects for the STM transport shall be supported by the STM-TC. In the case of multiple bearer channels, the corresponding STM-TCs shall be represented by independent sets of indicators. The anomalies and defect indicators shall comply with clause L.1.

11.3.2.3 Anomalies and defects for PTM transport

The anomalies and defects for the PTM transport shall be supported by the PTM-TC. In the case of multiple bearer channels, the corresponding PTM-TCs shall be represented by independent sets of indicators. The anomalies and defect indicators shall comply with clause L.3.

11.3.3 Power-related primitives

11.3.3.1 Near-end primitives

Loss of power (*lpr*): This primitive occurs when the VTU power supply (mains) voltage drops below the manufacturer-determined level required for proper VTU operation. An *lpr* terminates when the power level exceeds the manufacturer-determined minimum power level.

11.3.3.2 Far-end primitives

Far-end loss of power (*flpr*): This primitive detected at the far end is reported by the *flpr* indicator, which shall be coded 1 to indicate that no *lpr* is being reported and shall be coded 0 for the next 3 *lpr* indicator transmissions to indicate that an *flpr* (i.e., "dying gasp") is being reported. An *flpr* occurs when 2 or more out of 3 consecutively received *lpr* indicators are set to ZERO. An *flpr* terminates when, for a period of 0.5 seconds, the received *lpr* indicator bit is set to ONE and no near-end *los* is present.

11.3.4 INM primitives

INM-related primitives represent anomalies related to PMD and PMS-TC sublayers.

11.3.4.1 INM INPEQ histogram primitives

If INM_INPEQ_FORMAT=0, then the INM INPEQ histogram shall be configured with a linear scale as follows:

- $INMAINPEQ_1..INMAINPEQ_{16}$: every $INMAINPEQ_i$ is a primitive detected at the near-end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.2.2.1) is exactly *i* DMT symbols;
- $INMAINPEQ_{17}$ is a primitive detected at the near-end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.2.2.1) is strictly more than 16 DMT symbols.

If INM_INPEQ_FORMAT=1 then the INM INPEQ histogram shall be configured with a logarithmic scale as follows ($\lfloor x \rfloor$ denotes rounding to the lower integer):

- $INMAINPEQ_1..INMAINPEQ_{16}$: every $INMAINPEQ_i$ is a primitive detected at the near-end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.2.2.1) falls in the range from $\lfloor 1.33^{i+1} \rfloor$ to $\lfloor 1.33^{i+2} \rfloor - 1$ DMT symbols, both boundaries inclusive;
- $INMAINPEQ_{17}$ is a primitive detected at the near-end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.2.2.1) is at least $\lfloor 1.33^{17+1} \rfloor = 169$ DMT symbols.

NOTE – The logarithmic scale gives rise to the following possible INMAINPEQ histogram ranges (in DMT symbols). It gives a finer granularity for the higher probability short duration impulses, whilst still capturing some information about the longer duration events.

1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**
1 – 1	2 – 2	3 – 3	4 – 4	5 – 6	7 – 8	9 – 12	13 – 16	17 – 22	23 – 29	30 – 39	40 – 53	54 – 71	72 – 94	95 – 126	127 – 168	≥ 169

11.3.4.2 INM total measurement primitive

- INMAME: is a primitive detected at the near-end only. This indication occurs every time a data symbol is processed by the impulse noise sensor (INS).

11.3.4.3 INM inter-arrival time histogram primitives

- $INMAIAT_0$ is a primitive detected at the near-end only. This anomaly occurs when the reported value of IAT falls in the range from 2 to INMIATO-1, both boundaries inclusive;
- $INMAIAT_1..INMAIAT_6$: every $INMAIAT_i$ is a primitive detected at the near-end only. This anomaly occurs when the reported value of IAT falls in the range from $(INMIATO + (i - 1)\times(2^{INMIATS}))$ to $(INMIATO - 1) + i\times(2^{INMIATS})$, both boundaries inclusive;
- $INMAIAT_7$ is a primitive detected at the near-end only. This anomaly occurs when the reported value of IAT falls in the range from $INMIATO + 6\times(2^{INMIATS})$ to infinity.

11.4 OAM parameters

The system may support and use the relevant OAM parameters for the VTU-O and VTU-R, as specified in clauses 7.2 and 7.3 of [ITU-T G.997.1]. Specifically, these are:

- Line-related and Path-related performance parameters;
- Line-related and Path-related configuration parameters; and
- Inventory parameters.

Test parameters shall be computed and formatted as specified in clause 11.4.1 to be reported in the format specified in [ITU-T G.997.1].

11.4.1 Test parameters

The test parameters are measured by the PMD transmit or receive function and shall be reported on request to the near-end VME. Test parameters can be used to identify possible issues with the physical loop and to check for adequate physical media performance margin at acceptance and after repair verification, or at any other time following the initialization of the VDSL2 system.

The following test parameters shall be passed on request from the receive PMD function to the near-end VME:

- Channel characteristics function H(*f*) per subcarrier (CCF-ps);
- Quiet line noise PSD QLN(*f*) per subcarrier (QLN-ps);
- Signal-to-noise Ratio SNR(*f*) per subcarrier (SNR-ps);
- Loop attenuation per band (LATN-pb);
- Signal attenuation per band (SATN-pb);
- Signal-to-noise ratio margin per band (SNRM-pb);
- Signal-to-noise ratio margin for the ROC (SNRM-ROC);
- Attainable net data rate (ATTNDR);
- Far-end actual aggregate transmit power (ACTATP);
- Far-end actual impulse noise protection (INP_act); and
- Far-end actual impulse noise protection of the ROC (INP_act-ROC).

The following test parameter shall be passed on request from the transmit PMD function to the near-end VME:

- Near-end actual aggregate transmit power (ACTATP).

The purposes of making the above information available are:

- H(*f*) can be used to analyse the physical copper loop condition;
- QLN(*f*) can be used to analyse the crosstalk;
- SNR(*f*) can be used to analyse time-dependent changes in crosstalk levels and loop attenuation (such as due to moisture and temperature variations); and
- The combination of H(*f*), QLN(*f*) and SNR(*f*) can be used to help determine why the data rate is not equal to the maximum data rate for a given loop.

The detailed diagnostic information H(*f*) and QLN(*f*) would be most useful during showtime. However, requesting this would place an undue computational burden on the VTUs. Thus, the combination of complete information on the channel (H(*f*) and QLN(*f*)) during initialization combined with initialization and showtime SNR(*f*) is provided as a reasonable compromise. This combination of data will allow greater analysis of the loop conditions than traditional methods and will reduce interruptions to both VDSL2 and the underlying service that traditional diagnostic methods require.

The quiet line noise (QLN), signal-to-noise ratio (SNR), and channel characteristics in format (Hlin, Hlog) shall be represented by subcarrier groups. The number of subcarriers, G, in one subcarrier group shall be equal to:

$$G = pow2\ (\Theta/512)$$

where the function $pow2(x)$ takes the nearest power of 2 greater than or equal to x and Θ is the highest subcarrier index of the transmitter SUPPORTEDCARRIERS set if the parameter is measured during the channel discovery phase; or the last subcarrier index of the transmitter MEDLEY set in other cases.

Specific carrier sets to be used during showtime and loop diagnostic mode are summarized in Table 11-42 (N/A indicates that a parameter is not applicable).

Table 11-42 – Value of *G* for various phases of operation

Test parameter	Normal operation	Loop diagnostic mode	
	Showtime	Channel discovery	Channel analysis and exchange
QLN	SUPPORTEDCARRIERS	SUPPORTEDCARRIERS	N/A
HLOG	SUPPORTEDCARRIERS	SUPPORTEDCARRIERS	N/A
HLIN	N/A	N/A	MEDLEY
SNR	MEDLEY	N/A	MEDLEY

Valid values of *G* are 1, 2, 4 and 8.

11.4.1.1 Definition of test parameters

11.4.1.1.1 Channel characteristics function per subcarrier group (CCF-ps)

The channel characteristics function H(*f*) is a quantity that is related to the values of the (complex) source and load impedances. A simplified definition is used in which the source and load impedances are the same and equal to a real value R_N. The channel characteristics function H(*f*) is associated with a two-port network, normalized to a chosen reference resistance R_N. H(*f*) shall be defined as a complex value, equal to the U2/U1 voltage ratio (see Figures 11-3 and 11-4).

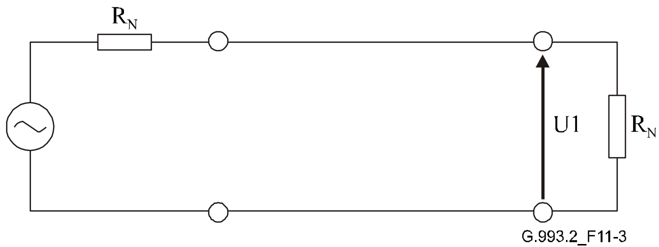

Figure 11-3 – Voltage across the load

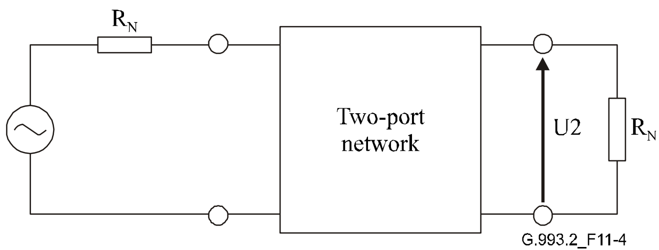

Figure 11-4 – Voltage across the load with a two-port network inserted

The measurement of a channel characteristics function is the result of the cascade of three functions:

1) The transmitter filter characteristics function;

2) The channel characteristics function; and

3) The receiver filter characteristics function.

NOTE – The channel characteristics function corresponds to the $H_{channel}(f)$ function used in the definition of the far-end crosstalk (see clause 7.4.1 of [ITU-T G.996.1]).

The objective is to provide means by which the channel characteristics can be accurately identified. Therefore, it is necessary for the receive PMD function to report an estimate of the channel characteristics. This task may prove to be a difficult one given the fact that the receive PMD function only observes the cascade of all three elements of the channel. The passband part of the reported H(*f*), which is most essential to debug possible issues with the physical loop, is not expected to significantly depend upon the receiver filter characteristics (not including receiver automatic gain control (AGC)). The receive PMD function shall therefore invert the gain (AGC) it has applied to the received signal and do a best effort attempt to remove the impact of the near-end receiver filter characteristics. The result is then a best estimate of how the receiver views the passband channel characteristics plus the transmitter filter characteristics. Because the in-band portion of the spectrum is also expected not to significantly depend upon the transmitter filter characteristics, this result is considered a sufficient estimate of the channel characteristics for desired loop conditioning applications.

Two formats for the channel characteristics are defined:

- Hlin(*f*): a format providing complex values on a linear scale; and
- Hlog(*f*): a format providing magnitude values on a base 10 logarithmic scale.

For Hlog(*f*), the receive PMD function shall also use the value of the PSD at the U interface of the transmit PMD function (as conveyed in messages during initialization) to remove the impact of the far-end transmit filter characteristics.

For Hlin(*f*), if the channel characteristics are reported over the VTU-O OAM interface (see Figure 5-3), the VTU-O shall do a best effort attempt to remove the impact of the near-end transmit filter characteristics from the channel characteristics measured at the VTU-R. If the channel characteristics are reported over the VTU-R OAM interface, the VTU-R shall do a best effort attempt to remove the impact of the near-end transmit filter characteristics from the channel characteristics measured at the VTU-O.

Hlin(*f*) shall be sent to the far-end VME during the loop diagnostic mode and shall be sent on request to the near-end VME during the loop diagnostic mode.

Hlog(*f*) shall be measured by the receive PMD function during the loop diagnostic mode and initialization. The measurement shall not be updated during showtime. Hlog(*f*) shall be sent to the far-end VME during the loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send Hlog(*f*) to the far-end VME on request during showtime.

In loop diagnostic mode, both Hlin(*f*) and Hlog(*f*) shall be measured, because the corrections that can be done, relative to the receiver and/or transmitter filter characteristics with Hlin(*f*) and Hlog(*f*), may differ.

Hlin(*f*) and Hlog(*f*) shall be measured over a 1 second time period in loop diagnostic mode. In initialization, the VTU should do its best to optimize the accuracy of the Hlog(*f*) measurement; however, it shall measure at least 256 symbols, and shall indicate the measurement period (in symbols, represented as a 16-bit unsigned value) to the far-end VME (see clause 11.2.3.11).

The channel characteristics function Hlin($k \times G \times \Delta f$) shall be the value of the channel characteristics on the subcarrier $i = k \times G$. It shall be represented in linear format by a scale factor and a normalized complex number $a(k) + j \times b(k)$, $k = 0$ to 511. The scale factor shall be coded as a 16-bit unsigned integer. Both $a(k)$ and $b(k)$ shall be coded as 16-bit twos complement signed integers. The value of Hlin($k \times G \times \Delta f$) shall be defined as:

$$\text{Hlin}(k \times G \times \Delta f) = (scale/2^{15}) \times (a(k) + j \times b(k))/2^{15}$$

In order to maximize precision, the scale factor, *scale*, shall be chosen such that $\max(|a(k)|, |b(k)|)$ over all k is equal to $2^{15} - 1$.

This data format supports an Hlin(*f*) granularity of 2^{-15} and an Hlin(*f*) dynamic range of approximately 96 dB (+6 dB to −90 dB). The portion of the scale factor range above 0 dB is necessary because, due to manufacturing variations in signal path gains and filter responses, short loops may appear to have a gain rather than a loss.

An Hlin($k \times G \times \Delta f$) value indicated as $a(k) = b(k) = -2^{15}$ is a special value. It indicates that:

- no measurement could be done for this subcarrier because it is out of the transmitter MEDLEY set or its $g_i = 0$; or
- the attenuation is out of the range to be represented.

The channel characteristics function Hlog($k \times G \times \Delta f$) shall be the magnitude of the channel characteristics at subcarrier $k \times G$. It shall be represented in base 10 logarithmic format by an integer number $m(k)$, where $k = 0$ to 511. The $m(k)$ shall be coded as 10-bit unsigned integers. The value of Hlog($k \times G \times \Delta f$) shall be defined as:

$$\text{Hlog}(k \times G \times \Delta f) = 6 - (m(k)/10)$$

This data format supports an Hlog(f) granularity of 0.1 dB and an Hlog(f) dynamic range of approximately 102 dB (+6 dB to −96 dB).

An Hlog($k \times G \times \Delta f$) value indicated as $m(k) = 2^{10} - 1$ is a special value. It indicates:

- that no measurement could be done for this subcarrier because it is out of the transmitter SUPPORTEDCARRIERS set; or
- that the attenuation is out of the range to be represented.

11.4.1.1.2 Quiet line noise PSD per subcarrier group (QLN-ps)

The quiet line noise PSD QLN(f) for a particular subcarrier is the rms level of the noise present on the loop when no VDSL2 signals are present on the loop. The received virtual noise PSD as defined in SNRM_MODE=2, SNRM_MODE=3, SNRM_MODE=4, and SNRM_MODE=5 shall not be taken into account in QLN(f).

The quiet line noise PSD QLN(f) per subcarrier shall be measured by the receive PMD function during loop diagnostic mode and initialization. The measurement shall not (i.e., cannot) be updated during showtime. The QLN(f) shall be sent to the far-end VME during loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the QLN(f) to the far-end VME on request during showtime.

The objective is to provide means by which the quiet line noise PSD can be accurately identified. Therefore, it would be necessary for the receive PMD function to report an estimate of the quiet line noise PSD. This task may prove to be a difficult one given the fact that the receive PMD function observes the noise through the receiver filter. The passband part of the reported QLN-ps, which is most essential to debug possible issues with the physical loop, is not expected to significantly depend upon the receiver filter characteristics (not including receiver AGC). The receive PMD function shall therefore invert the gain (AGC) it has applied to the received signal and do a best effort attempt to remove the impact of the near-end receiver filter characteristics. The result is then a best estimate of how the receiver views the passband quiet line noise PSD. This result is considered a sufficient estimate of the quiet line noise PSD for desired loop conditioning applications.

The receive PMD function shall measure the QLN(f) in a time interval when no VDSL2 signals are present on the loop (i.e., the near-end and far-end transmitters are inactive). The quiet line noise PSD QLN($i \times \Delta f$) shall be measured over a 1 second time interval in loop diagnostic mode. In initialization, the VTU should do its best to optimize the accuracy of the QLN(f) measurement, however, it shall measure over at least 256 symbols, and shall indicate the measurement period (in symbols, represented as a 16-bit unsigned value) to the far-end VME (see clause 11.2.3.11).

The quiet line noise PSD QLN($k \times G \times \Delta f$) shall be the average of the power values of quiet line noise on the subcarriers $k \times G$ to $((k+1) \times G) - 1$. It shall be represented as an 8-bit unsigned integer $n(k)$, where $k = 0$ to 511. The value of QLN($k \times G \times \Delta f$) shall be defined as QLN($k \times G \times \Delta f$) = $-23 - (n(k)/2)$ dBm/Hz. This data format supports a QLN(f) granularity of 0.5 dB with a range of values for QLN(f) from −150 to −23 dBm/Hz.

A QLN($k \times G \times \Delta$f) value indicated as $n(k)$ = 255 is a special value. It indicates that:

- no measurement could be done for this subcarrier group because one of its subcarriers is out of the transmitter SUPPORTEDCARRIERS set; or
- the quiet line noise PSD is out of the range to be represented.

11.4.1.1.3 Signal-to-noise ratio per subcarrier group (SNR-ps)

The signal-to-noise ratio SNR(f) for a particular subcarrier is a real value that shall represent the ratio between the received signal power and the received noise power for that subcarrier. The received virtual noise PSD as defined in SNRM_MODE=2, SNRM_MODE=3, and SNRM_MODE=4 shall not be taken into account in SNR(f).

The signal-to-noise ratio SNR(f) per subcarrier shall be measured by the receive PMD function in loop diagnostic mode and initialization. The measurement may be updated autonomously and shall be updated on request during showtime. The SNR(f) shall be sent to the far-end VME during loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the SNR(f) to the far-end VME on request during showtime.

The receive PMD function shall measure the signal-to-noise ratio SNR(f) with the transmit PMD function in a MEDLEY or showtime state. The signal-to-noise ratio SNR(f) shall be measured over a 1 second time interval in loop diagnostic mode. In initialization and showtime, the VTU should do its best to minimize the SNR(f) measurement time, however it shall measure over at least 256 symbols, and it shall indicate the measurement period (in symbols, represented as a 16-bit unsigned value) to the far-end VME (see clause 11.2.3.11).

The signal-to-noise ratio SNR($k \times G \times \Delta f$) shall be the average of the base 10 logarithmic value of the signal-to-noise ratio on the subcarriers $k \times G$ to $((k+1) \times G) - 1$. It shall be represented as an 8-bit unsigned integer $snr(k)$, where k = 0 to 511. The value of SNR($k \times G \times \Delta f$) shall be defined as SNR($k \times G \times \Delta f$) = −32 + ($snr(k)$/2) dB. This data format supports an SNR($k \times G \times \Delta f$) granularity of 0.5 dB and an SNR($k \times G \times \Delta f$) dynamic range of 127 dB (−32 to 95 dB).

An SNR($k \times G \times \Delta f$) value indicated as $snr(k)$ = 255 is a special value. It indicates that:

- no measurement could be done for this subcarrier group because one of its subcarriers is out of the transmitter MEDLEY set or its $g_i = 0$; or
- the signal-to-noise ratio is out of the range to be represented.

11.4.1.1.4 Loop attenuation per band (LATN-pb)

The loop attenuation in the m^{th} downstream band is denoted as LATN_D(m), and the loop attenuation in the m^{th} upstream band is denoted as LATN_U(m). For ease of notation, this clause provides requirements and definitions in terms of the downstream loop attenuation, but the same definitions and requirements also apply to LATN_U(m).

The loop attenuation of the m^{th} downstream band (LATN_D(m)) is the squared magnitude of the channel characteristics function H(f) (as defined in clause 11.4.1.1.1) averaged over all subcarriers of the m^{th} downstream band, converted to dB. LATN_D(m) shall be defined as:

$$\text{LATN\_D}(m) = -10 \times \log_{10}\left(\frac{\sum_{i=n1}^{n2} \left| \text{H}(i \times \Delta f) \right|^2}{N\_D(m)} \right)$$

with $N\_D(m)$ (the number of subcarriers in the m^{th} downstream band) = $n2-n1+1$ where $n1$ and $n2$ are the indices of the first and the last subcarriers of this band, respectively, and H(f) is represented by Hlin(f) in loop diagnostic mode and by Hlog(f) in initialization (with conversion of $\log_{10}$ to linear values for use in the above equation).

If one or more H(f) values could not be measured because they are out of the transmitter SUPPORTEDCARRIERS set (see clause 11.4.1.1.1), then the LATN_D(m) shall be calculated as an average of H(f) values over the number of subcarriers for which valid values of H(f) are available.

The loop attenuation shall be calculated by the receive PMD function during loop diagnostic mode and initialization. The calculation shall not be updated during showtime. The loop attenuation shall be sent to the far-end VME during initialization and loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the LATN to the far-end VME on request during showtime.

The loop attenuation per downstream band LATN_D(m) shall be represented as a 10-bit unsigned integer *latn*, with the value of LATN_D(m) defined as LATN_D(m) = *latn*/10 dB. This data format supports a LATN_D(m) granularity of 0.1 dB and an LATN_D(m) dynamic range of 102.2 dB (0 to 102.2 dB).

A LATN_D(m) value indicated as *latn* = 1023 is a special value. It indicates that the loop attenuation is out of the range that can be represented.

11.4.1.1.5 Signal attenuation per band (SATN-pb)

The signal attenuation in the m^{th} downstream band is denoted as SATN_D(m), and the signal attenuation in the m^{th} upstream band is denoted as SATN_U(m). For ease of notation, this clause provides requirements and definitions in terms of the downstream signal attenuation, but the same definitions and requirements also apply to SATN_U(m).

The signal attenuation of the m^{th} downstream band, SATN_D(m), is defined as the difference in dB between the power received at the near-end and that transmitted from the far end in the m^{th} downstream band.

Mathematically, this corresponds to:

$$\text{SATN\_D}(m) = \text{TXpower\_dBm\_D}(m) - \text{RXpower\_dBm\_D}(m)$$

During initialization and loop diagnostic mode, the received signal power in dBm, RXpower_dBm_D(m), shall be computed as the received subcarrier power, summed over those subcarriers of this band that are in the MEDLEYds set. During transmission of O-P-MEDLEY and R-P-MEDLEY, the transmit PSD for subcarriers in the MEDLEYds set is at the MREFPSDds level. Therefore, the received signal power shall be fine-tuned with the g_i values for each subcarrier in the MEDLEYds set to estimate the signal power that will be received during showtime.

Mathematically, this corresponds to:

$$\text{RXpower\_dBm\_D(m)} = 10 \times \log_{10}\left(\sum_{i \in (MEDLEYds \cap DS(m))} \left(\text{Received\_subcarrier\_power\_mW}(i) \times g_i^2 \right) \right)$$

During showtime, the received signal power in dBm, Rxpower_dBm_D(m), shall be computed as the received subcarrier power in showtime, summed over those subcarriers of this band that are in the MEDLEYds set.

Mathematically, this corresponds to:

$$\text{RXpower\_dBm\_D(m)} = 10 \times \log_{10}\left(\sum_{i \in (MEDLEYds \cap DS(m))} \left(\text{Received\_subcarrier\_power\_mW}(i) \right) \right)$$

In both equations, MEDLEYds ∩ DS(m) denotes all subcarriers of the MEDLEYds set that fall into the m^{th} downstream band, Received_subcarrier_power_mW is the received power on subcarrier i expressed in milli-Watts, and g_i is the gain (linear scale) for subcarrier i.

The received power for SATN_U(m) shall be computed in the same way, but using subcarriers from the MEDLEYus set falling into the m^{th} upstream band.

For the SATN value determined during initialization, the received signal power for each subcarrier i in the MEDLEYds set shall be fine-tuned with the g_i value conveyed in the O-PMD (for the upstream direction) and R-PMD (for the downstream direction) messages to estimate the signal power that will be received during showtime. During loop diagnostic mode, the fine tuning shall be restricted to using g_i values 0 (for subcarriers to which no bits can be allocated) and 1 (for subcarriers to which at least one bit can be allocated). For the SATN value determined during Showtime, the received signal subcarrier power shall be taken as measured.

The transmitted signal power in dBm, TXpower_dBm_D(m), corresponds to the part of the NOMATP (see clause 10.3.4.2.1) falling in this band. It shall be computed as the aggregate transmit power, summed over the subcarriers of this band that are in the MEDLEYds set. During transmission of O-P-MEDLEY, the transmit PSD for subcarriers in the MEDLEYds set is at the MREFPSDds level. Therefore, the transmitted signal power shall be fine-tuned with the g_i values for each subcarrier in the MEDLEYds set to estimate the signal power that will be transmitted during showtime.

Mathematically, this corresponds to:

$$\text{TXpower\_dBm\_D}(m) = 10 \times \log_{10} \Delta f + 10 \times \log_{10}\left(\sum_{i \in \text{MEDLEYds} \cap \text{DS}(m))} \left(10^{\frac{\text{MREFPSD}[i]}{10}} \times g_i^2 \right) \right)$$

where MEDLEYds $\cap$ DS(m) denotes all subcarriers of the MEDLEYds set that fall into the m^{th} downstream band, *MREFPSD[i]* is the value of MREFPSDds for subcarrier i in dBm/Hz as conveyed by the O-PRM message, g_i is the gain (linear scale) for subcarrier i, and Δf is the subcarrier spacing in Hz.

The transmit power for SATN_U(m) shall be computed in the same way, but using subcarriers from the MEDLEYus set falling into the m^{th} upstream band, and the value of MREFPSD[i] is the value of MREFPSDus for subcarrier i in dBm/Hz as conveyed by the R-PRM message.

For the SATN value determined during initialization, the transmit signal power for each subcarrier i in the MEDLEYds set shall be fine-tuned with the g_i value conveyed in the O-PMD (for the upstream direction) and R-PMD (for the downstream direction) messages to estimate the transmit signal power during showtime. During loop diagnostic mode, the fine tuning shall be restricted to using g_i values 0 (for subcarriers to which no bits can be allocated) and 1 (for subcarriers to which at least one bit can be allocated). For the SATN value determined during showtime, the transmitted signal power shall be fine-tuned with the active g_i values for each subcarrier in the MEDLEY set.

The signal attenuation shall be measured by the receive PMD function during loop diagnostic mode and initialization (i.e., estimate the signal attenuation at the start of showtime). The measurement may be updated autonomously and shall be updated on request during showtime. The signal attenuation shall be sent to the far-end VME during initialization and loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the SATN_D(m) to the far-end VME on request during showtime.

The signal attenuation per downstream band, SATN_D(m), shall be represented as a 10-bit unsigned integer *satn*, with the value of SATN_D(m) defined as SATN_D(m) = *satn*/10 dB. This data format supports an SATN_D(m) granularity of 0.1 dB and an SATN_D(m) dynamic range of 102.2 dB (0 to 102.2 dB).

An SATN_D(m) value indicated as *satn* = 1023 is a special value. It indicates that the signal attenuation is out of range to be represented.

11.4.1.1.6 Signal-to-noise ratio margin

11.4.1.1.6.1 General definition of signal-to-noise ratio margin

The signal-to-noise ratio margin is the maximum increase (scalar gain, in dB) of the reference noise PSD (at all relevant frequencies), such that the BER of each bearer channel does not exceed 10^{-7} (see clause 9.8), without any change of PMD parameters (e.g., bits and gains) and PMS-TC parameters (e.g., L_p, FEC parameters). The BER is referenced to the output of the PMS-TC function (i.e., the α/β interface).

The definition of the reference noise PSD depends on the control parameter SNRM_MODE.

11.4.1.1.6.1.1 SNRM_MODE = 1

SNRM_MODE = 1 is a mandatory capability for both VTUs.

The reference noise PSD equals the received current-condition external noise PSD only, as measured by the near-end transceiver (i.e., equal to the PSD of the noise measured by the near-end transceiver at the constellation decoder or other relevant internal reference point when the only noise source is the external stationary noise applied to the U interface and no internal noise sources are present).

NOTE – Mathematically this can be illustrated by:

$$\text{Received\_External\_Noise\_PSD} = |\, H_{\text{RXfilter}}(f)\, |^2 \times \text{External\_Noise\_PSD\_at\_U\_interface}$$

11.4.1.1.6.1.2 SNRM_MODE = 2

SNRM_MODE = 2 is an optional capability for both VTUs.

The reference noise PSD equals the maximum of the received current-condition external noise PSD (as defined in SNRM_MODE=1) and the received virtual noise PSD, at a common internal reference point.

The received virtual noise PSD shall be determined by the transceiver as defined in the following equation.

$$\text{Received\_Virtual\_Noise\_PSD} = |\, H(f)\, |^2 \times \text{TXREFVN}$$

where TXREFVN is the transmitter-referred virtual noise PSD MIB parameter.

$|H(f)|^2$ is calculated as:

$$|H(f)|^2 = \frac{\text{Actual\_Received\_Signal\_PSD}}{\text{Actual\_Transmit\_Signal\_PSD}}$$

where:

- Actual_Transmit_Signal_PSD is the actual transmit signal PSD at the far-end transmitter as calculated by the near-end transceiver.
- Actual_Received_Signal_PSD is the actual received signal PSD at the near-end transceiver as measured by the near-end transceiver (i.e., equal to the PSD measured by the near-end transceiver at the constellation decoder or other relevant internal reference point) during initialization and Showtime.

Mathematically this can be expressed as:

$$\text{Actual\_Received\_Signal\_PSD} = |H_{\text{RXfilter}}(f)|^2 \times \text{Received\_Signal\_PSD\_at\_U\_interface}$$

NOTE – A measurement of the current-condition external noise PSD could be overly optimistic, as it only represents a snapshot in time, not taking into account the future increase in noise PSD (e.g., due to additional VDSL2 lines being switched on). The SNRM_MODE=2 is defined to prevent the VTU's bit loading algorithm from assigning an overly optimistic number of bits to a subcarrier. This is achieved by defining (via the transmitter-referred virtual noise PSD parameter TXREFVN) an anticipated noise PSD, which may be a function of frequency that can be used for bit loading.

This method can be used to avoid or reduce periods with excessive BER and retrains, in order to assure service quality and stability. It is expected that the configuration, via the MIB, is based on anticipated service penetration and noise environment.

11.4.1.1.6.1.3 SNRM_MODE = 3

SNRM_MODE = 3 is an optional capability for upstream (i.e., VTU-O receiver) only.

The reference noise PSD equals the maximum of the received current-condition external noise PSD (as defined in SNRM_MODE=1) and the received virtual noise PSD, at a common internal reference point.

The received virtual noise PSD shall be determined by the transceiver as defined in the following equation.

$$\text{Received\_Virtual\_Noise\_PSD} = |H_{RXfilter}(f)|^2 \times \text{RXREFVN}$$

where:

$|H_{RXfilter}(f)|$	is the magnitude of the transfer function between the U-interface and the common internal reference point, and
RXREFVN	is the receiver-referred virtual noise PSD MIB parameter.

NOTE – A measurement of the current-condition external noise PSD could be overly optimistic, as it only represents a snapshot in time, not taking into account the future increase in noise PSD (e.g., due to additional VDSL2 lines being switched on). The SNRM_MODE = 3 is defined to prevent the VTU's bit loading algorithm from assigning an overly optimistic number of bits to a subcarrier. This is achieved by defining (via the receiver-referred virtual noise PSD parameter RXREFVN) an anticipated noise PSD, which may be a function of frequency that can be used for bit loading.

This method can be used to avoid or reduce periods with excessive BER and retrains, in order to assure service quality and stability. It is expected that the configuration, via the MIB, is based on anticipated service penetration and noise environment.

11.4.1.1.6.1.4 SNRM_MODE = 4

SNRM_MODE = 4 is an optional capability for both upstream and downstream.

The reference noise PSD equals the maximum of the received current-condition external noise PSD (as defined in SNRM_MODE=1) and the received virtual noise PSD, at a common internal reference point.

In the downstream direction, the received virtual noise PSD shall be determined by the transceiver as defined in the following equation:

$$\text{Received\_Virtual\_Noise\_PSD} = |H(f)|^2 \times \text{TXREFVNds}$$

where:

$|H(f)|^2$	is calculated as in SNRM_MODE=2, and
TXREFVNds	is the transmitter-referred virtual noise PSD control parameter as conveyed in O-SIGNATURE.

In SNRM_MODE= 4, the control parameter TXREFVNds is a combination of the configuration parameter TXREFVNds and the configuration parameter TXREFVNSFds as described in clause 11.4.2.5.

In the upstream direction, the received virtual noise PSD shall be determined by the transceiver as defined in the following equation:

$$\text{Received\_Virtual\_Noise\_PSD} = |H_{RXfilter}(f)|^2 \times \text{RXREFVNus}$$

where:

$\lvert H_{RXfilter}(f)\rvert$	is the magnitude of the transfer function between the U-interface and the common internal reference point, and
RXREFVNus	is the receiver-referred virtual noise PSD control parameter.

In SNRM_MODE=4, the control parameter RXREFVNus is a combination of the configuration parameter RXREFVNus and the configuration parameter RXREFVNSFus as described in clause 11.4.2.4.

NOTE 1 – A measurement of the current-condition external noise PSD could be overly optimistic, as it only represents a snapshot in time, not taking into account the future increase in noise PSD (e.g., due to additional VDSL2 lines being switched on). SNRM_MODE = 4 is defined to prevent the VTU's bit loading algorithm from assigning an overly optimistic number of bits to a subcarrier. This is achieved by defining (via the virtual noise PSD parameter and the virtual noise scaling factor parameter) an anticipated noise PSD, which may be a function of frequency that can be used for bit loading. This method can be used to avoid or reduce periods with excessive BER and retrains, in order to assure service quality and stability. It is expected that the configuration, via the MIB, is based on anticipated service penetration and noise environment.

NOTE 2 – SNRM_MODE=4 is similar to SNRM_MODE=3 in the upstream direction and SNRM_MODE=2 in the downstream direction. Per line configuration in SNRM_MODE=2, 3 and 4 may be accomplished by configuration of a different value for the configuration parameter Virtual Noise (VNus or VNds, see clauses 7.3.1.7.3 and 7.3.1.7.4 of [ITU-T G.997.1]) on a line by line basis. Per line configuration in SNRM_MODE=4 may also be accomplished by the configuration of a common value for the configuration parameter Virtual Noise (VNus or VNds) for a group of lines combined with the configuration of a scaling factor (TXREFVNSFds or RXREFVNSFus) on a line by line basis.

11.4.1.1.6.1.5 SNRM_MODE = 5

SNRM_MODE = 5 is an optional capability for both upstream and downstream. A transceiver supporting operation with SNRM_MODE = 5 shall also support SRA.

The SNRM_MODE = 5 can be applied in both directions (i.e., for both VTU-O and VTU-R receivers). If SNRM_MODE = 5 is applied in at least one direction, then an *msg* bit rate of at least 64 kbit/s shall be used in both directions.

The reference noise PSD is equal to the maximum of the received current-condition external noise PSD (as defined in SNRM_MODE=1) and the received showtime adaptive virtual noise (SAVN) PSD, at a common internal reference point.

The received SAVN PSD value shall be determined by the transceiver as defined in the following equation.

$$\text{Received\_SAVN\_PSD} = \lvert H(f)\rvert^2 \times \text{TXREFSAVN}(f)$$

where TXREFSAVN*(f)* is the transmitter-referred SAVN PSD parameter at frequency *f*, and $\lvert H(f)\rvert^2$ is calculated as defined in clause 11.4.1.1.6.1.2.

The value of TXREFSAVN is set at initialization and may be updated during showtime. At initialization, the downstream SAVN PSD and upstream SAVN PSD are determined by the VTU-O. The initial value of the downstream SAVN PSD (TXREFSAVNds) is communicated to the VTU-R in the TXREFVNds field in the O-SIGNATURE message (see clause 12.3.3.2.1.1). During showtime, the VTU-O may update the value of the downstream SAVN PSD based on noise estimations or actual noise measurements at the VTU-R, e.g., noise changes due to other lines going in and out of the low power mode (see Annex E of [ITU-T G.998.4]). The update of TXREFSAVNds is communicated to the VTU-R via a SAVN-Update eoc command (see clause 11.2.3.16). During showtime, the VTU-O may also update the value of the upstream SAVN PSD based on noise estimations or actual noise measurements at the near-end. The VTU-O may estimate the SNR at the remote end by using the Test Parameter Read command defined in clause 11.2.3.11. Other estimation methods are for further study. Upon an update of the received SAVN PSD, the VTU shall verify whether the update requires a modification of the bit loading. If a modification is required, the VTU shall perform an SRA to adjust the bit loading in accordance with the updated value of the received SAVN PSD.

NOTE 1 – To identify changes in the SNR, the VTU-O ME may request the SNR read right after implementation of an SRA command, or a bit swap command, or after SOS.

The range of TXREFSAVNds, TXREFSAVNus values shall be in the boundaries determined by the CO-MIB (parameters TXREFSAVNds-MIN/MAX, TXREFSAVNus-MIN/MAX defined in Table 11-42.1).

NOTE 2 – To set the initial values of TXREFSAVNds, TXREFSAVNus, the VTU-O may take in account the values TXREFVNds and TXREFVNus (see clause 11.4.2.1) provided by the CO-MIB, respectively.

Table 11-42.1 – CO-MIB configuration parameters related to SAVN

Configuration parameter	Definition
TXREFSAVNds-MIN	TXREFSAVNds-MIN indicates a minimum value, such that TXREFSAVNds is greater than or equal to TXREFSAVNds-MIN for all applicable subcarriers. The valid range is –150 to –23 dBm/Hz in steps of 0.5dB.
TXREFSAVNds-MAX	TXREFSAVNds-MAX indicates a maximum value, such that TXREFSAVNds is less than or equal to TXREFSAVNds-MAX for all applicable subcarriers. The valid range is –150 to –23 dBm/Hz in steps of 0.5dB.
TXREFSAVNus-MIN	TXREFSAVNus-MIN indicates a minimum value, such that TXREFSAVNus is greater than or equal to TXREFSAVNus-MIN for all applicable subcarriers. The valid range is –150 to –23 dBm/Hz in steps of 0.5dB.
TXREFSAVNus-MAX	TXREFSAVNus-MAX indicates a maximum value, such that TXREFSAVNus is less than or equal to TXREFSAVNus-MAX for all applicable subcarriers. The valid range is –150 to –23 dBm/Hz in steps of 0.5dB.

11.4.1.1.6.2 Signal-to-noise ratio margin parameter (SNRM)

The signal-to-noise ratio margin parameter, SNRM, is the signal-to-noise ratio margin (as defined in clause 11.4.1.1.6.1) measured over all subcarriers, except the subcarriers assigned to the ROC, in a transmission direction for which $b_i > 0$. The received virtual noise PSD as defined in clauses 11.4.1.1.6.1.2, 11.4.1.1.6.1.3, 11.4.1.1.6.1.4, and 11.4.1.1.6.1.5 shall be taken into account, respectively, when configured in SNRM_MODE=2, SNRM_MODE=3, SNRM_MODE=4, and SNRM_MODE=5.

The signal-to-noise ratio margin shall be measured by the receive PMD function during initialization. The measurement may be updated autonomously and shall be updated on request during showtime. The SNRM shall be sent to the far-end VTU during initialization and loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the SNRM to the far-end VME on request during showtime.

To determine the SNRM, the receive PMD function must be able to first determine the bits and gains table. During loop diagnostic mode, the receive PMD function shall use the special value to indicate that the SNRM value was not measured.

The signal-to-noise ratio margin in the downstream direction shall be represented as a 10-bit twos complement signed integer *snrm*, with the value of SNRMds defined as SNRMds = *snrm*/10 dB. This data format supports an SNRMds granularity of 0.1 dB and an SNRMds dynamic range of 102.2 dB (−51.1 to +51.1 dB).

An SNRMds value indicated as *snrm* = −512 is a special value. It indicates that the signal-to-noise ratio margin is out of the range to be represented. During loop diagnostic mode, the special value shall be used to indicate that the SNRMds value was not measured.

The same definition and representation shall apply to the signal-to-noise ratio margin in the upstream direction, SNRMus.

11.4.1.1.6.3 Signal-to-noise ratio margin per band (SNRM-pb)

The signal-to-noise ratio margin in the m^{th} downstream band is denoted as SNRM_D(*m*), and the signal-to-noise ratio margin in the m^{th} upstream band is denoted as SNRM_U(*m*). For ease of notation, this clause provides requirements and definitions in terms of the downstream signal-to-noise ratio margin, but the same definitions and requirements also apply to SNRM_U(*m*).

The signal-to-noise ratio margin per band parameter SNRM-pb is the signal-to-noise ratio margin (as defined in clause 11.4.1.1.6.1) measured over all subcarriers in a particular band for which $b_i > 0$. The received virtual noise PSD as defined in clauses 11.4.1.1.6.1.2, 11.4.1.1.6.1.3, 11.4.1.1.6.1.4, and 11.4.1.1.6.1.5 shall be taken into account, respectively, when configured in SNRM_MODE=2, SNRM_MODE=3, SNRM_MODE=4, and SNRM_MODE=5.

The signal-to-noise ratio margin per band is the maximum increase (in dB) in the received noise power that can be tolerated in this band, such that the VTU can still meet all target BERs over all bearer channels.

The signal-to-noise ratio margin per band shall be measured by the receive PMD function during initialization. The measurement may be updated autonomously and shall be updated on request during showtime. The signal-to-noise ratio margin per band shall be sent to the far-end VME during initialization and loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the SNRM-pb to the far-end VME on request during showtime.

To determine the SNRM-pb, the receive PMD function must be able to first determine the bits and gains table. During loop diagnostic mode, the receive PMD function shall use the special value to indicate that the SNRM-pb value was not measured.

The signal-to-noise ratio margin per downstream band shall be represented as a 10-bit twos complement signed integer *snrm*, with the value of SNRM_D(*m*) defined as SNRM_D(*m*) = *snrm*/10 dB. This data format supports an SNRM_D(*m*) granularity of 0.1 dB and an SNRM_D(*m*) dynamic range of 102.2 dB (−51.1 to +51.1 dB).

An SNRM_D(*m*) value indicated as *snrm* = −512 is a special value. It indicates that the signal-to-noise ratio margin is out of the range to be represented. During loop diagnostic mode, the special value shall be used to indicate that the SNRM_D(*m*) value was not measured.

11.4.1.1.6.4 Signal-to-noise ratio margin for the ROC (SNRM-ROC)

The SNRM-ROC is the signal-to-noise ratio margin related to transmission of the ROC, as defined in clause 9.1, Figure 9-2. The definition of SNRM-ROC is as in clause 11.4.1.1.6.1 applied to the MPS-TC and BER=10^{-7} for all bits transmitted over latency path #0, Figure 9-2. The SNRM-ROC shall be measured over all subcarriers assigned to the ROC for which $b_i > 0$ in a transmission direction. The received virtual noise PSD as defined in clause 11.4.1.1.6.1.2 shall be taken into account when configured in SNRM_MODE=2.

The SNRM-ROC shall be measured by the receive PMD function during initialization. The measurement may be updated autonomously and shall be updated on request during Showtime. The SNRM-ROC shall be sent to the far-end VTU during initialization and shall be sent on request to the near-end VME at any time. The near-end VME shall send the SNRM-ROC to the far-end VME on request during Showtime. The receive PMD function shall use a special value to indicate that the SNRM value was not measured (e.g., in Loop Diagnostic mode or if the ROC is not enabled or not supported).

The SNRM-ROC shall use the same representation as defined for SNRM in clause 11.4.1.1.6.2.

11.4.1.1.7 Attainable net data rate (ATTNDR)

The attainable net data rate (ATTNDR) is the maximum net data rate that the receive PMS-TC and PMD functions are designed to support, assuming certain conditions. Assumed conditions depend on the control parameter attndr_method (single method applicable to both upstream and downstream directions): either the basic method (*attndr_method*=0, see clause 11.4.1.1.7.1) or the improved methods (*attndr_method*=1 or 2, see clause 11.4.1.1.7.2). The *attndr_method* value is sent to the VTU-R during initialization (see clause 12.3.5.2.2.1). The actual *attndr_method* value is reported through the CO-MIB parameter ATTNDR_ACTMETHOD (see clause 7.5.1.41.1 of [ITU-T G.997.1]).

The improved Attainable net data rate method shall be used if the CO-MIB parameter ATTNDR_METHOD (see clause 7.3.1.15.1 of [ITU-T G.997.1]) is set to a value of 1 or 2, and the mode is supported by the VTU-O and by the VTU-R. In this case the value of the control parameter *attndr_method* shall have the same value as the CO-MIB configuration parameter ATTNDR_METHOD (i.e., *attndr_method*=ATTNDR_METHOD). Otherwise, the basic attainable net data rate method shall be used. In this case the value of the control parameter *attndr_method* is zero (i.e., *attndr_method*=0).

To accurately determine the attainable net data rate (ATTNDR), the receive PMD function must be able to first determine the bits and gains table. Therefore, during loop diagnostic mode, the ATTNDR value for upstream and downstream shall be calculated as:

$$\text{ATTNDR} = \sum_{i=0}^{MSI} \min\left\{\text{round}\left[\log_2\left(1+10^{(SNR(i\times\Delta f)-\text{SNRGAP}-\text{TARSNRM})/10}\right)\right], 15\right\} \times 4\,\text{kbit/s}$$

with SNR($i \times \Delta f$) in dB as defined in clause 11.4.1.1.3, but accounting for the received virtual noise PSD when configured in SNRM_MODE=2, SNRM_MODE=3, SNRM_MODE=4, or SNRM_MODE=5 and SNRGAP= 9.75 dB (Note 1).

NOTE 1 – The SNRGAP value is defined for a 10^{-7} bit error ratio on 4-QAM (no coding gain, $INP\_min_0 = 0$).

NOTE 2 – The value calculated for ATTNDR during loop diagnostic mode may not be identical to the value calculated during Showtime with the same PMD parameters and under the same loop conditions.

The attainable net data rate shall be calculated by the receive PMS-TC and PMD functions during loop diagnostic mode and initialization. The measurement may be updated autonomously and shall be updated on request during showtime. The attainable net data rate shall be sent to the far-end VME during loop diagnostic mode and shall be sent on request to the near-end VME at any time. The near-end VME shall send the ATTNDR to the far-end VME on request during showtime.

The attainable net data rate shall be represented as a 32-bit unsigned integer *attndr*, with the value of ATTNDR defined as ATTNDR = *attndr* bit/second. This data format supports an ATTNDR granularity of 1 bit/s.

No special value is defined.

11.4.1.1.7.1 The basic attainable net data rate method

Support of the basic attainable net data rate method is mandatory.

The attainable net data rate is the maximum net data rate that the receive PMS-TC and PMD functions are designed to support, assuming the following conditions:

- single bearer channel and single latency operation;
- target SNR margin equal to the configured TARSNRMds/TARSNRMus downstream and upstream, respectively;

- latency not to exceed the highest latency configured for the bearer channel (*delay_act*$_0$ ≤ *delay_max*$_0$);
- accounting for all coding gains available (e.g., trellis coding, FEC) within the latency bound;
- accounting for the channel characteristics at the instant of measurement; and
- accounting for the received virtual noise PSD when configured in SNRM_MODE=2, SNRM_MODE=3, SNRM_MODE=4, or SNRM_MODE=5.

NOTE – The conditions of the basic method in this version of the Recommendation are unchanged compared to the conditions for ATTNDR calculation in previous versions of this Recommendation. This set of conditions did not specify a number of conditions to calculate ATTNDR, which lead to vendor discretionary behaviour in the reported ATTNDR values. It is recommended that implementations that are upgraded from a previous version of this Recommendation to this version of this Recommendation, in case of *attndr_method*=0 use the same vendor discretionary behaviour as used when operating according to the previous version of this Recommendation.

11.4.1.1.7.2 The improved attainable net data rate method

Support of the improved attainable net data rate method is optional.

The attainable net data rate is the maximum net data rate that the receive PMS-TC and PMD functions are designed to support, assuming the conditions of the basic attainable net data rate method (see clause 11.4.1.1.7.1) and the following conditions:

- if operating in single bearer channel mode, the ATTNDR calculation shall use single bearer channel mode. If operating in dual bearer channel mode, the ATTNDR calculation shall use single bearer channel mode (see clause 8.1.3.1);
- if operating in single latency mode, the ATTNDR calculation shall use single latency mode. If operating in dual latency mode, the ATTNDR calculation shall use single latency mode. If operating in single latency with ROC mode, the ATTNDR calculation shall use single latency with ROC mode (see clause 9.1);
- if the control parameter *attndr_method* is set to a value of 1, the VTU-O and VTU-R shall use the impulse noise protection limit *INP_min*$_0$ with value as indicated in O-TPS (see clause 12.3.5.2.1.2). If O-TPS configures for dual bearer channel mode, then the highest *INP_min*$_n$ over the bearer channels shall be used;
- if the control parameter *attndr_method* is set to a value of 2, the VTU-O and VTU-R shall use an impulse noise protection limit *INP_min*$_0$ = 0;
- the VTU-O and VTU-R shall use the maximum delay *delay_max*$_0$ with value as indicated in O-TPS (see clause 12.3.5.2.1.2). If O-TPS configures for dual bearer channel mode, then the lowest *delay_max*$_n$ over the bearer channels shall be used;
- use of Erasure decoding or not is identical to usage on the bearer channels;
- taking into account the framing limitations;
- if ATM bonding is set in the ITU-T G.994.1 bonding code tree, the VTU-O and VTU-R shall use the minimum delay *delay_min*$_0$ as defined in clause L.2.7. If O-TPS configures for dual bearer channel mode, then the lowest *delay_min*$_n$ over the bearer channels shall be used;
- Taking into account the value of the ATTNDR_MDOSPLITparameter;
- Net data rate is not limited by the configured maximum net data rate;
- Channel initialization policy CIP = 0;
- Transmit PSD is equal to MREFPSD for all subcarriers for which $g_i \neq 0$.

  NOTE 1 – The ATTNDR value may be lower due to possible transmit power reductions, as a consequence of configured MAXMARGIN setting, configured MAXNDR setting and vendor discretionary transmit power reductions (e.g., subcarriers with $g_i = 0$, due to AFE dynamic range, …).

NOTE 2 – The basic method does not specify a number of conditions to calculate ATTNDR, which leads to vendor discretionary behaviour in the reported ATTNDR values. The improved method defines additional conditions to reduce variation of reported ATTNDR values over implementations.

When the ATTNDR value is reported during Showtime, the following parameters used in the calculation of the ATTNDR shall also be reported over the eoc with the ATTNDR value (see Table 11-30), and also be reported in the CO-MIB with the ATTNDR value (see clauses 7.5.1.19 and 7.5.1.20 of [ITU-T G.997.1]):

- $ATTNDR\_INP\_act_0$ (see clauses 7.5.1.41.2 and 7.5.41.3 of [ITU-T G.997.1]);
- $ATTNDR\_delay\_act_0$ (see clauses 7.5.1.41.6 and 7.5.41.7 of [ITU-T G.997.1]).

The parameter $ATTNDR\_INP\_act_0$ is the far-end actual impulse noise protection used in the calculation of the ATTNDR. The actual impulse noise protection $ATTNDR\_INP\_act_0$ shall be represented as an 8-bit unsigned integer $attndr\_inp\_act_0$, with the value of $ATTNDR\_INP\_act_0$ defined as $ATTNDR\_INP\_act_0 = attndr\_inp\_act_0 / 10$ DMT symbols. This data format supports an $ATTNDR\_INP\_act_0$ granularity of 0.1 DMT symbol. The range is from 0 DMT symbols (represented as 0) to 25.4 DMT symbols (represented as 254). The value 255 is a special value indicating an $ATTNDR\_INP\_act_0$ higher than 25.4 DMT symbols.

The parameter $ATTNDR\_delay\_act_0$ is the far-end actual delay used in the calculation of the ATTNDR. The actual delay $ATTNDR\_delay\_act_0$ shall be represented as an 8-bit unsigned integer $attndr\_delay\_act_0$, with the value of $ATTNDR\_delay\_act_0$ defined as $ATTNDR\_delay\_act_0 = attndr\_delay\_act_0 / 10$ ms. This data format supports an $ATTNDR\_delay\_act_0$ granularity of 0.1 ms. The range is from 0 ms (represented as 0) to 25.4 ms (represented as 254). The value 255 is a special value indicating an $ATTNDR\_DELAY\_act_0$ higher than 25.4 ms.

11.4.1.1.8 Actual aggregate transmit power (ACTATP)

The actual aggregate transmit power (ACTATP) is the total amount of output power delivered by the transmit PMD function to the U reference point at tip-and-ring (in dB), at the instant of measurement. The transmit PMD function shall take the NOMATP (see clause 10.3.4.2.1) as a best estimate of the near-end actual aggregate transmit power.

The receive PMD function shall take NOMATP (see clause 10.3.4.2.1) as a best estimate of the far-end actual aggregate transmit power.

The near-end and far-end actual aggregate transmit power shall be calculated by the VTU during initialization using the assigned values of g_i. The measurement may be updated autonomously and shall be updated on request during showtime. The near-end and far-end actual aggregate transmit power shall be sent on request to the near-end VME. The near-end VME shall send the near-end and far-end ACTATP to the far-end VME on request during showtime.

To determine the near-end ACTATP, the transmit PMD function must first receive the bits and gains table from the receive PMD function. Therefore, during loop diagnostic mode, the g_i values shall be determined as value 1 (for all subcarriers in the MEDLEY set).

The actual aggregate transmit power shall be represented as a 10-bit twos complement signed integer *actatp*, with the value of ACTATP defined as ACTATP = *actatp*/10 dBm. This data format supports an ACTATP granularity of 0.1 dB, with an ACTATP dynamic range of 62 dB (−31 to + 31 dBm).

An ACTATP value indicated as *actatp* = −512 is a special value. It indicates that the actual aggregate transmit power is out of the range to be represented.

11.4.1.1.9 Actual impulse noise protection (INP_act)

The actual impulse noise protection INP_act_n of bearer channel #n is defined in clause 9.6.

The actual impulse noise protection INP_act_n of bearer channel #n shall be represented as an 8-bit unsigned integer inp_act_n, with the value of INP_act_n defined as $\text{INP\_act}_n = inp\_act_n / 10$ DMT symbols. This data format supports an INP_act_n granularity of 0.1 DMT symbol. The range is from 0 DMT symbols (represented as 0) to 25.4 DMT symbols (represented as 254). The value 255 is a special value indicating an INP_act_n higher than 25.4 DMT symbols.

11.4.1.1.10 Actual impulse noise protection of the ROC (INP_act-ROC)

The INP_act-ROC is the actual impulse noise protection of the ROC, as defined in clause 9.1, Figure 9-2. It shall be computed as $INP\_act\text{-}ROC = INP\_no\_erasure_0$ (see clause 9.6).

The parameter INP_act-ROC is expressed in DMT symbols.

The parameter ACTINP-ROC in the CO-MIB is expressed in multiples of T_{4k}. The value is coded in fractions of T_{4k} periods with a granularity of 0.1 of a period. The range is from 0 to 25.4. At the Q-reference point, a special value indicates an ACTINP-ROC higher than 25.4. The T_{4k} is defined in clause 10.4.4.

The VTU-O shall calculate ACTINP-ROC as follows:

- For 4.3125 kHz subcarrier spacing:

  ACTINP-ROC = MIN($\lfloor 10 \times$ INP_act-ROC $\rfloor$ / 10; 25.5)

- For 8.625 kHz subcarrier spacing:

  ACTINP-ROC = MIN($\lfloor 5 \times$ INP_act-ROC $\rfloor$ / 10; 25.5)

11.4.1.2 Accuracy of test parameters

This clause defines accuracy requirements for test parameters defined in clause 11.4.1.1. The accuracy requirement is expressed as a tolerance relative to a reference value. Both the reference value and the allowed tolerance are defined in this clause.

The accuracy requirements of test parameters are optional. A VTU may comply with the accuracy requirements for all or a subset of the test parameters. For subcarriers and subcarrier groups to which accuracy requirements do not apply (see clauses 11.4.1.2.1 to 11.4.1.2.8), the VTU shall report test parameters as defined in clause 11.4.1.1 but the accuracy of these reported values may not comply with the accuracy requirements further defined in this clause.

NOTE – The measurement of test parameter reference values involves the use of test equipment. The accuracy requirements defined in this clause do not take into account test equipment tolerance. Test equipment tolerance is out of the scope of this Recommendation and is to be added to the tolerances defined in this clause.

11.4.1.2.1 Accuracy of channel characteristics function per subcarrier group (CCF-ps)

11.4.1.2.1.1 Accuracy of Hlog($k \times G \times \Delta f$)

The downstream HLOG(f) reference value for frequency $k \times G \times \Delta f$ shall be defined as:

$$HLOG\_reference\_ds(k \times G \times \Delta f) = MREFPSDds(k \times G \times \Delta f) - PSD\_UR2(k \times G \times \Delta f),$$

where PSD_UR2($k \times G \times \Delta f$) is the PSD measured at the U-R2 reference point with the VTU-O connected to the loop and frozen in the O-P-MEDLEY stage of initialization with the SOC in the O-IDLE state, and with the VTU-R replaced by an R_N=100 Ohm resistance terminating the loop.

The upstream HLOG(f) reference value for frequency $k \times G \times \Delta f$ shall be defined *as*:

$$HLOG\_reference\_us(k \times G \times \Delta f) = MREFPSDus(k \times G \times \Delta f) - PSD\_UO2(k \times G \times \Delta f),$$

where PSD_UO2($k \times G \times \Delta f$) is the PSD measured at the U-O2 reference point with the VTU-R connected to the loop and frozen in the R-P-MEDLEY stage of initialization with the SOC in the R-IDLE state, and with the VTU-O replaced by an R_N=100 Ohm resistance terminating the loop.

NOTE 1 – The feature to freeze a VTU in the MEDLEY stage of initialization exists solely to allow a test bed to be constructed for the purpose of measuring the Hlog(f) reference values. It applies only to specific transceivers serving as the 'transmit transceiver' of the test environment and is not a requirement for compliance with this Recommendation.

The receiving VTU shall measure the Hlog(f) values under the same loop, noise, temperature, and configuration settings that are used for measuring the Hlog(f) reference values.

The accuracy requirements for the Hlog($k \times G \times \Delta f$) shall only apply to those subcarrier groups with an SNR (as defined in clause 11.4.1.1.3) ≥ 12 dB, where the SNR is the SNR value measured during initialization, after the channel discovery phase.

The accuracy requirements for the downstream Hlog($k \times G \times \Delta f$) (ITU-T G.997.1 parameter HLOGpsds):

- shall not apply to subcarrier groups that contain subcarriers from the downstream BLACKOUT set;
- shall not apply to subcarrier groups that contain subcarriers in the RFI bands or that contain any of the 15 subcarriers adjacent to each side of the RFI bands, and
- shall only apply to subcarrier groups for which all subcarriers within the group fall within the following frequency ranges (defined as a part of the passband):
  - For Annex A, Masks D-32, D-48, and D-64 of Table A.8:
    - Subcarrier 92 to Subcarrier 869 and Subcarrier 1206 to Subcarrier 1971 for profiles 8a, 8b, 8c, 8d, 12a, 12b and 17a.
  - For Annex A Mask D-128 of Table A.8:
    - Subcarrier 184 to Subcarrier 869 and Subcarrier 1206 to Subcarrier 1971 for profiles 8a, 8b, 8c, 8d, 12a, 12b and 17a.
  - For Annex B band plan 998 of Table B.1:
    - Subcarrier 92 to Subcarrier 869 and Subcarrier 1206 to Subcarrier 1971 for profiles 8a, 8b, 8c, 8d, 12a, 12b and 17a.
  - For Annex B band plan 997 of Table B.1:
    - Subcarrier 92 to Subcarrier 695 and Subcarrier 1183 to Subcarrier 1634 for profiles 8a, 8b, 8c, 8d, 12a, 12b and 17a.
  - For Annex C, Masks in Tables C.1, C.2, C.5 and C.6:
    - Subcarrier 92 to Subcarrier 869 and Subcarrier 1206 to Subcarrier 1971 for profiles 8a, 8b, 8c, 8d, 12a, 12b and 17a.
  - For Annex C, Masks in Table C.9:
    - Subcarrier 214 to Subcarrier 869 and Subcarrier 1206 to Subcarrier 1971 for profiles 8a, 8b, 8c, 8d, 12a, 12b and 17a.

Accuracy requirements for Annex B band plans 998ADE are for further study.

Accuracy requirements for Profile 30a are for further study.

Accuracy requirements for Annex N band plans are for further study.

Accuracy requirements outside these specified ranges are for further study.

The accuracy requirements for the upstream Hlog($k \times G \times \Delta f$) (ITU-T G.997.1 parameter HLOGpsus):

- shall not apply to subcarrier groups that contain subcarriers from the upstream BLACKOUT set;
- shall not apply to subcarrier groups that contain subcarriers in the RFI bands or that contain any of the 15 subcarriers adjacent to each side of the RFI bands; and

- shall only apply to subcarrier groups for which all subcarriers within the group fall within the following frequency ranges (defined as a part of the passband):
  - For Annex A, Annex B Band Plan 998 of Table B.1 and Annex C:
    - Subcarrier 870 to Subcarrier 1205 for profiles 8a, 8b, 8c and 8d.
    - Subcarrier 870 to Subcarrier 1205 and Subcarrier 1972 to Subcarrier 2782 for profiles 12a, 12b and 17a.
  - For Annex B Band Plan 997 of Table B.1:
    - Subcarrier 696 to Subcarrier 1182 for profile 8c.
    - Subcarrier 696 to Subcarrier 1182 and Subcarrier 1635 to Subcarrier 2047 for profiles 8a, 8b and 8d.
    - Subcarrier 696 to Subcarrier 1182 and Subcarrier 1635 to Subcarrier 2782 for profiles 12a, 12b and 17a.

Accuracy requirements for Annex B band plans 998ADE are for further study.

Accuracy requirements for Profile 30a are for further study.

Accuracy requirements for Annex N band plans are for further study.

Accuracy requirements for the US0 band (for all relevant profiles) are for further study.

Accuracy requirements outside these frequency ranges are for further study.

NOTE 2 – Having such specified ranges for accuracy requirements avoids variations due to the tolerances and effects of the filtering at the low ends of the passband and of the effects of folding at the high end of the passband.

The accuracy requirements for downstream and upstream Hlog($k \times G \times \Delta f$) shall only apply for those subcarrier groups where the loop impedance (Z_{loop}) falls within the following ranges for all the subcarriers in the group:

- Impedance magnitude is between 100 Ohm and 120 Ohm;
- Impedance imaginary component is between –20 Ohm and 0 Ohm.

Z_{loop} is defined as the impedance seen by the receiving transceiver under test, looking into the loop, including the transmitting transceiver connected to the loop at the far end.

Accuracy requirements for downstream and upstream Hlog($k \times G \times \Delta f$), for frequencies where Z_{loop} falls outside this range, are for further study.

NOTE 3 – Appendix II provides an informative discussion of the effects on the accuracy of Hlog(f) measurements caused by impedance mismatch between a nominal 100 Ohm termination of the loop and possible termination impedances (Z_{VTU}) actually provided by a VTU.

For each subcarrier group where the accuracy requirement for downstream Hlog($k \times G \times \Delta f$) applies (based on its subcarrier indexes and downstream SNR($k \times G \times \Delta f$) value only, and not considering restrictions related to its Z_{loop} values), and where HLOGps_reference_ds($k \times G \times \Delta f$) is above –90 dB, a downstream Hlog($k \times G \times \Delta f$) value different from the special value defined in clause 11.4.1.1.1 shall be reported.

For each subcarrier group where the accuracy requirement for downstream Hlog($k \times G \times \Delta f$) applies, and where HLOGps_reference_ds($k \times G \times \Delta f$) is above –90 dB, the absolute error between the downstream Hlog($k \times G \times \Delta f$) and HLOGps_reference_ds($k \times G \times \Delta f$) shall be ≤ 3 dB.

The accuracy requirements for downstream Hlog($k \times G \times \Delta f$) shall apply to its measurement either during Initialization or in Loop Diagnostic mode.

For each subcarrier group where the accuracy requirement for upstream Hlog($k \times G \times \Delta f$) applies (based on its subcarrier indexes and upstream SNR($k \times G \times \Delta f$) value only, and not considering restrictions related to its Z_{loop} values), and where HLOGps_reference_us($k \times G \times \Delta f$) is above –90 dB, an upstream Hlog($k \times G \times \Delta f$) value different from the special value defined in clause 11.4.1.1.1 shall be reported.

For each subcarrier group where the accuracy requirement for upstream Hlog($k \times G \times \Delta f$) applies, and where HLOGps_reference_us($k \times G \times \Delta f$) is above –90 dB, the absolute error between the upstream Hlog($k \times G \times \Delta f$) and the HLOGps_reference_us($k \times G \times \Delta f$) shall be ≤ 3 dB.

The accuracy requirements for upstream Hlog($k \times G \times \Delta f$) shall apply to its measurement either during Initialization or in Loop Diagnostic mode.

11.4.1.2.1.2 Accuracy of Hlin($k \times G \times \Delta f$)

The accuracy requirements for the magnitude of Hlin($k \times G \times \Delta f$) are the same as those for Hlog($k \times G \times \Delta f$) in clause 11.4.1.2.1.1.

There is no accuracy requirement for the phase of Hlin($k \times G \times \Delta f$).

11.4.1.2.2 Accuracy of quiet line noise PSD per subcarrier group (QLN-ps)

The downstream QLN(f) reference value for subcarrier group k including subcarriers $i = k \times G$ to $((k + 1) \times G) - 1$ shall be defined as:

$$\text{QLNps\_reference\_ds}(k \times G \times \Delta f) = \frac{1}{G} \sum_{i=kG}^{(k+1)G-1} \text{PSDps\_UR2}(i \times \Delta f)$$

where PSDps_UR2($i \times \Delta f$) is the PSD (in logarithmic scale) at frequency $i \times \Delta f$ measured at the U-R2 reference point in the downstream bands, after initialization of the line up to an O-P-QUIET stage, in which stage the VTU-O is frozen and the VTU-R subsequently replaced by an R_N=100 Ohm resistance.

The upstream QLN(f) reference value for subcarrier group k including subcarriers $i = k \times G$ to $((k + 1) \times G) - 1$ shall be defined as:

$$\text{QLNps\_reference\_us}(k \times G \times \Delta f) = \frac{1}{G} \sum_{i=kG}^{(k+1)G-1} \text{PSDps\_UO2}(i \times \Delta f)$$

where PSDps_UO2($i \times \Delta f$) is the PSD (in logarithmic scale) at frequency $i \times \Delta f$ measured at the U-O2 reference point in the upstream bands, after initialization of the line up to an R-P-QUIET stage, in which stage the VTU-R is frozen and the VTU-O subsequently replaced by an R_N=100 Ohm resistance.

NOTE 1 – The feature to freeze a VTU in a QUIET stage exists solely to allow a test bed to be constructed for the purpose of measuring the QLN(f) reference value. It applies only to specific transceivers serving as the 'transmit transceiver' of the test environment and is not a requirement for compliance with this Recommendation.

The receiving VTU shall measure the QLN(f) values under the same loop, noise, temperature, and configuration settings as are used for measuring the QLN(f) reference values.

The accuracy requirements for the downstream QLN($k \times G \times \Delta f$) (ITU-T G.997.1 parameter QLNpsds) shall apply to the subcarrier groups in the same frequency bands and with the same loop impedance (Z_{loop}) restrictions as where the downstream Hlog($k \times G \times \Delta f$) accuracy requirements apply (see clause 11.4.1.1.2).

The accuracy requirements for the upstream QLN($k \times G \times \Delta f$) (ITU-T G.997.1 parameter QLNpsus) shall apply to the subcarrier groups in the same frequency bands and with the same loop impedance (Z_{loop}) restrictions as where the upstream Hlog($k \times G \times \Delta f$) accuracy requirements apply (see clause 11.4.1.1.2).

Accuracy requirements outside these frequency ranges are for further study.

NOTE 2 – Having such specified ranges for accuracy requirements avoids variations due to the tolerances and effects of the filtering at the low ends of the passband and of the effects of folding at the high end of the passband.

For each subcarrier group where the accuracy requirement for downstream QLN($k \times G \times \Delta f$) applies (based on its subcarrier indexes only, and not considering restrictions related to its Z_{loop} values), and where QLNps_reference_ds($k \times G \times \Delta f$) is above –130 dBm/Hz, a downstream QLN($k \times G \times \Delta f$) value different from the special value defined in clause 11.4.1.1.2 shall be reported.

For each subcarrier group where the accuracy requirement for downstream QLN($k \times G \times \Delta f$) applies, and where QLNps_reference_ds($k \times G \times \Delta f$) is above –130 dBm/Hz, the absolute error between the downstream QLN($k \times G \times \Delta f$) and the QLNps_reference_ds($k \times G \times \Delta f$) shall be ≤ 3.0 dB. To account for sinusoidal noise sources internal to the VTU-R, this requirement does not apply to up to 5 clusters of N consecutive subcarrier groups per 2.2 MHz bandwidth, which can be selected at the VTU-R vendor's discretion, with N = 1 + ceil(W/G) and W = 12.

For each subcarrier group where the accuracy requirement for downstream QLN($k \times G \times \Delta f$) applies, and where QLNps_reference_ds($k \times G \times \Delta f$) is above –130 dBm/Hz, the sample variance of downstream QLN($k \times G \times \Delta f$) measurements (within a 10 minute measurement window, and under the same loop, noise, temperature, and configuration settings) shall be ≤ 0.5 dB. To account for sinusoidal noise sources internal to the VTU-R, this requirement does not apply to up to 5 clusters of N consecutive subcarrier groups per 2.2 MHz bandwidth, which can be selected at the VTU-R vendor's discretion, with N = 1 + ceil(W/G) and W = 12.

The accuracy requirements for downstream QLN($k \times G \times \Delta f$) shall apply to its measurement either during Initialization or in Loop Diagnostic mode.

For each subcarrier group where the accuracy requirement for upstream QLN($k \times G \times \Delta f$) applies (based on its subcarrier indexes only, and not considering restrictions related to its Z_{loop} values), and where QLNps_reference_us($k \times G \times \Delta f$) is above –120 dBm/Hz, an upstream QLN($k \times G \times \Delta f$) value different from the special value defined in clause 11.4.1.1.2 shall be reported.

For each subcarrier group where the accuracy requirement for upstream QLN($k \times G \times \Delta f$) applies, and where QLNps_reference_us($k \times G \times \Delta f$) is above –120 dBm/Hz, the absolute error between the upstream QLN($k \times G \times \Delta f$) and the QLNps_reference_us($k \times G \times \Delta f$) shall be ≤ 3.0 dB. To account for sinusoidal noise sources internal to the VTU-O, this requirement does not apply to up to 10 clusters of N consecutive subcarrier groups per 2.2 MHz bandwidth, which can be selected at the VTU-O vendor's discretion, with N = 1 + ceil(W/G) and W = 12.

For each subcarrier group where the accuracy requirement for upstream QLN($k \times G \times \Delta f$) applies, and where QLNps_reference_us($k \times G \times \Delta f$) is above –120 dBm/Hz, the statistical sample variance of upstream QLN($k \times G \times \Delta f$) measurements (within a 10 minute measurement window, and under the same loop, noise, temperature, and configuration settings) shall be ≤ 0.5 dB. To account for sinusoidal noise sources internal to the VTU-O, this requirement does not apply to up to 10 clusters of N consecutive subcarrier groups per 2.2 MHz bandwidth, which can be selected at the VTU-O vendor's discretion, with N = 1 + ceil(W/G) and W = 12.

The accuracy requirements for upstream QLN($k \times G \times \Delta f$) shall apply to its measurement either during Initialization or in Loop Diagnostic mode.

11.4.1.2.3 Accuracy of signal-to-noise ratio per subcarrier group (SNR-ps)

Noise PSD changes over time shall be reflected in the reported SNR($k \times G \times \Delta f$). This clause defines accuracy requirements for the change in SNR($k \times G \times \Delta f$) over a time interval [T1,T2], relative to a reference value. The downstream and upstream reference values for subcarrier group k including subcarriers $i = k \times G$ to $((k + 1) \times G) - 1$ are defined as:

$$\Delta \text{SNRps\_reference\_ds}(k \times G \times \Delta f) = \text{Noise\_PSDps\_UR2\_T1}(k \times G \times \Delta f) - \text{Noise\_PSDps\_UR2\_T2}(k \times G \times \Delta f)$$

ΔSNRps_reference_us($k \times G \times \Delta f$) = Noise_PSDps_UO2_T1($k \times G \times \Delta f$) – Noise_PSDps_UO2_T2($k \times G \times \Delta f$)

where:

- Noise_PSDps_UR2_T1 is the stationary noise PSD (in dBm/Hz) present on the line at the U-R2 reference point at time instant T1, and for at least one minute before T1;
- Noise_PSDps_UR2_T2 is the stationary noise PSD (in dBm/Hz) present on the line at the U-R2 reference point at time instant T2, and for at least one minute before T2;
- Noise_PSDps_UO2_T1 is the stationary noise PSD (in dBm/Hz) present on the line at the U-O2 reference point at time instant T1, and for at least one minute before T1;
- Noise_PSDps_UO2_T2 is the stationary noise PSD (in dBm/Hz) present on the line at the U-O2 reference point at time instant T2, and for at least one minute before T2.

These four Noise_PSDps's shall be measured by the same method as is used to measure the QLNps_reference (see clause 11.4.1.1.2) and before the SNR measurements. Before the actual measurements of SNR, the two noise PSDs (for time T1 and T2) shall be measured while the transmitting VTU is frozen in a QUIET state. Then the transmitting VTU is allowed to enter SHOWTIME and the SNR measurements are made under the same two Noise_PSDps's. The SNR measurements shall be made under the same loop and temperature conditions as the Noise_PSDps measurements.

The accuracy requirements for the downstream SNR($k \times G \times \Delta f$) (ITU-T G.997.1 parameter SNRpsds) shall apply to those subcarrier groups in the downstream passband where all of the following conditions hold:

- Subcarriers in the subcarrier group are at least 50 kHz away from the lower and higher passband edge;
- bi_T1(i) > 0 and bi_T2(i) > 0 for at least one subcarrier i in the subcarrier group (i between $k \times G$ and $(k+1) \times G - 1$ for subcarrier group k);
- Noise_PSDps_UR2_T1($k \times G \times \Delta f$) and Noise_PSDps_UR2_T2($k \times G \times \Delta f$) are larger than –110 dBm/Hz;
- (SNRps_T1($k \times G \times \Delta f$) – $\frac{1}{G}\sum_{i=kG}^{(k+1)G-1}$ gi_T1(i)) and (SNRps_T2($k \times G \times \Delta f$) – $\frac{1}{G}\sum_{i=kG}^{(k+1)G-1}$ gi_T2(i)) are both smaller than 40 dB,

where:

- gi_T1(i) and gi_T2(i) are the downstream fine gains (in dB) at time instants T1 and T2;
- bi_T1(i) and bi_T2(i) are the downstream bit loading at time instants T1 and T2;
- SNRps_T1($k \times G \times \Delta f$) and SNRps_T2($k \times G \times \Delta f$) are the downstream SNRs (in dB), measured during showtime, at time instants T1 and T2.

The accuracy requirements for the upstream SNR($k \times G \times \Delta f$) (ITU-T G.997.1 parameter SNRpsus) shall apply to those subcarrier groups in the upstream passband where all of the following conditions hold:

- Subcarriers in the subcarrier group are at least 50kHz away from the lower and higher passband edge;
- bi_T1(i) > 0 and bi_T2(i) > 0, for at least one subcarrier i in the subcarrier group (i between $k \times$ G and $(k+1) \times$ G – 1 for subcarrier group k);
- Noise_PSDps_UO2_T1($k \times G \times \Delta f$) and Noise_PSDps_UO2_T2($k \times G \times \Delta f$) are larger than –120 dBm/Hz;

- (SNRps_T1($k \times G \times \Delta f$) – $\frac{1}{G}\sum_{i=kG}^{(k+1)G-1}$ gi_T1(i)) and (SNRps_T2($k \times G \times \Delta f$) – $\frac{1}{G}\sum_{i=kG}^{(k+1)G-1}$ gi_T2(i)) are both smaller than 40 dB,

where:

- gi_T1(i) and gi_T2(i) are the upstream fine gains (in dB) at time instants T1 and T2;
- bi_T1(i) and bi_T2(i) are the upstream bit loading at time instants T1 and T2;
- SNRps_T1($k \times G \times \Delta f$) and SNRps_T2($k \times G \times \Delta f$) are the upstream SNRs (in dB), measured during showtime, at time instants T1 and T2.

If the line does not re-initialize over a time period T1 to T2, the following requirements shall be met for downstream subcarrier groups where the SNR($k \times G \times \Delta f$) accuracy requirement applies:

$$|(\text{SNRps\_T2}(k \times G \times \Delta f) - \frac{1}{G}\sum_{i=kG}^{(k+1)G-1} \text{gi\_T2}(i)) - (\text{SNRps\_T1}(k \times G \times \Delta f) - \frac{1}{G}\sum_{i=kG}^{(k+1)G-1} \text{gi\_T1}(i)) - \Delta\text{SNRps\_reference\_ds}(k \times G \times \Delta f)| \le 0.8 \text{ dB.}$$

Accuracy requirements for downstream subcarrier groups where (SNRps_T1 – gi_T1) or (SNRps_T2 – gi_T2) is greater than 40 dB, are for further study.

For each downstream subcarrier group where the SNR($k \times G \times \Delta f$) accuracy requirement applies, the sample variance of SNR($k \times G \times \Delta f$) measurements (expressed in dB and all samples taken within a 10 minute time interval, without line re-initialization in this time interval, and under the same loop, noise, temperature, and configuration settings) shall be equal to or smaller than 0.5, calculated as follows:

$$SNRps\_variance \le 0.5$$

where

$$SNRps\_variance = \frac{1}{N}\left(\sum_{i=1}^{N}\left(SNRps(i) - SNRps\_avg\right)^2\right)$$

$$SNRps\_avg = \frac{1}{N}\left(\sum_{i=1}^{N} SNRps(i)\right)$$

If the line does not re-initialize over a time period T1 to T2, the following requirements shall be met for upstream subcarrier groups where the SNR($k \times G \times \Delta f$) accuracy requirement applies:

$$|(\text{SNRps\_T2}(k \times G \times \Delta f) - \frac{1}{G}\sum_{i=kG}^{(k+1)G-1} \text{gi\_T2}(i)) - (\text{SNRps\_T1}(k \times G \times \Delta f) - \frac{1}{G}\sum_{i=kG}^{(k+1)G-1} \text{gi\_T1}(i)) - \Delta\text{SNRps\_reference\_us}(k \times G \times \Delta f)| \le 0.8 \text{ dB.}$$

Accuracy requirements for upstream subcarrier groups where (SNRps_T1 – gi_T1) or (SNRps_T2 – gi_T2) is greater than 40 dB, are for further study.

For each upstream subcarrier group where the SNR($k \times G \times \Delta f$) accuracy requirement applies, the sample variance of SNR($k \times G \times \Delta f$) measurements (expressed in dB and all samples taken within a 10 minute interval, without line re-initialization in this time interval, and under the same loop, noise, temperature, and configuration settings) shall be equal to or smaller than 0.5, calculated as follows:

$$SNRps\_variance \le 0.5$$

where

$$SNRps\_variance = \frac{1}{N}\left(\sum_{i=1}^{N}\left(SNRps(i) - SNRps\_avg\right)^2\right)$$

$$SNRps\_avg = \frac{1}{N}\left(\sum_{i=1}^{N} SNRps(i)\right)$$

NOTE – In verification tests, noise changes should be applied gradually over time, and not simultaneously at the U-O2 and U-R2 reference point, as not to force a re-initialization of the line.

11.4.1.2.4 Accuracy of loop attenuation per band (*LATN-pb*)

The downstream LATN reference value shall be defined as follows:

$$\text{LATN\_reference\_ds}(m) = -10 \times \log_{10}\left(\frac{\sum_{k=n1}^{n2}\left|\text{H\_reference\_ds}(k \times \Delta f)\right|^2}{N\_D(m)}\right)$$

In the equation above, $N\_D(m)$ (the number of subcarriers in the mth downstream band) = $n2-n1+1$ where $n1$ and $n2$ are the indices of the first and the last subcarriers of this band, respectively, and H_reference_ds($k \times \Delta f$) is defined as:

$$|\text{H\_reference\_ds}(k \times \Delta f)|^2 = 10^{\text{HLOG\_reference\_ds}(k \times \Delta f)/10} = 10^{(\text{MREFPSDds}(k \times \Delta f) - \text{PSD\_UR2}(k \times \Delta f))/10},$$

where PSD_UR2($k \times \Delta f$) is the PSD measured at the U-R2 reference point with the VTU-O connected to the loop and frozen in the O-P-MEDLEY stage of initialization with the SOC in the O-IDLE state, and with the VTU-R replaced by an R_N=100 Ohm resistance terminating the loop.

If one or more H_reference_ds(f) values could not be measured because they are out of the transmitter SUPPORTEDCARRIERS set (see clause 11.4.1.1.1), then the LATN_D(m) shall be calculated as an average of H(f) values over the number of subcarriers for which valid values of H(f) are available. Also,

The upstream LATN reference value shall be defined as follows:

$$\text{LATN\_reference\_us}(m) = -10 \times \log_{10}\left(\frac{\sum_{k=n1}^{n2}\left|\text{H\_reference\_us}(k \times \Delta f)\right|^2}{N\_U(m)}\right)$$

with $N\_U(m)$ (the number of subcarriers in the m^{th} upstream band) = $n2-n1+1$ where $n1$ and $n2$ are the indices of the first and the last subcarriers of this band, respectively. The value of H_reference_ds($k \times \Delta f$) is defined as:

$$|\text{H\_reference\_us}(k \times \Delta f)|^2 = 10^{\text{HLOG\_reference\_us}(k \times \Delta f)/10} = 10^{(\text{MREFPSDus}(k \times \Delta f) - \text{PSD\_UO2}(k \times \Delta f))/10}$$

where PSD_UO2($k \times \Delta f$) is the PSD measured at the U-O2 reference point with the VTU-R connected to the loop and frozen in the R-P-MEDLEY stage of initialization with the SOC in the R-IDLE state, and with the VTU-O replaced by an R_N=100 Ohm resistance terminating the loop.

If one or more H_reference_us(f) values could not be measured because they are out of the transmitter SUPPORTEDCARRIERS set (see clause 11.4.1.1.1), then the LATN_U(m) shall be calculated as an average of H(f) values over the number of subcarriers for which valid values of H(f) are available.

NOTE – The feature to freeze a VTU in the MEDLEY stage of initialization exists solely to allow a test bed to be constructed for the purpose of measuring the H(*f*) reference values. It applies only to specific transceivers serving as the 'transmit transceiver' of the test environment and is not a requirement for compliance with this Recommendation.

The receiving VTU shall measure the LATN values under the same loop, noise, temperature, and configuration settings that are used for measuring the LATN reference values.

For all downstream bands, *m*, that are active during loop diagnostics or initialization, the absolute error between LATN_D(*m*) and LATN_reference_ds(*m*) shall be equal to or smaller than 3 dB.

The accuracy requirements for downstream LATN_D(*m*) shall apply to its measurement either during Initialization or in Loop Diagnostic mode.

For all upstream bands, *m*, that are active during loop diagnostics or initialization, the absolute error between LATN_U(*m*) and LATN_reference_us(*m*) shall be equal to or smaller than 3 dB.

The accuracy requirements for upstream LATN_U(*m*) shall apply to its measurement either during Initialization or in Loop Diagnostic mode.

11.4.1.2.5 Accuracy of signal attenuation (*SATN*)

The downstream SATN reference value shall be defined as follows:

$$\text{SATN\_reference\_ds}(m) = \text{TXpower\_dBm\_reference\_ds}\ (m) - \text{RXpower\_dBm\_reference\_ds}\ (m)$$

The TXpower_dBm_reference_ds(*m*) is defined as:

$$\text{TXpower\_dBm\_reference\_ds}(m) = 10 \times \log_{10} \Delta f + 10 \times \log_{10}\left(\sum_{i \in \text{MEDLEYds} \cap \text{DS}(m)} \left(10^{\frac{\text{MREFPSD}[i]}{10}} \times g_i^2 \right) \right)$$

where MEDLEYds $\cap$ DS(*m*) denotes all subcarriers of the MEDLEYds set that fall into the *m*th downstream band, MREFPSD[*i*] is the value of MREFPSDds for subcarrier *i* in dBm/Hz, g_i is as defined in clause 11.4.1.1.5, and Δf is the subcarrier spacing in Hz.

The RXpower_dBm_reference_ds (*m*) is defined as:

$$\text{RXpower\_dBm\_reference\_ds}(m) = 10 \times \log_{10}\left(\sum_{i \in \text{MEDLEYds} \cap \text{DS}(m)} \left(10^{\frac{\text{PSD\_UR2}(i \times \Delta f)}{10}} \times g_i^2 \right) \right)$$

where PSD_UR2($i \times \Delta f$) is the PSD measured at the U-R2 reference point with the VTU-O connected to the loop and frozen in the O-P-MEDLEY stage of initialization with the SOC in the O-IDLE state, and with the VTU-R replaced by an R_N=100 Ohm resistance terminating the loop.

The upstream SATN reference value shall be defined as follows:

$$\text{SATN\_reference\_us}(m) = \text{TXpower\_dBm\_reference\_us}\ (m) - \text{RXpower\_dBm\_reference\_us}\ (m)$$

The TXpower_dBm_reference_us(*m*) is defined as:

$$\text{TXpower\_dBm\_reference\_us}(m) = 10 \times \log_{10} \Delta f + 10 \times \log_{10}\left(\sum_{i \in \text{MEDLEYus} \cap \text{US}(m)} \left(10^{\frac{\text{MREFPSD}[i]}{10}} \times g_i^2 \right) \right)$$

where MEDLEYus $\cap$ US(*m*) denotes all subcarriers of the MEDLEYus set that fall into the *m*th upstream band, MREFPSD[*i*] is the value of MREFPSDus for subcarrier *i* in dBm/Hz, g_i is as defined in clause 11.4.1.1.5, and Δf is the subcarrier spacing in Hz.

The RXpower_dBm_reference_us (*m*) is defined as:

$$\text{RXpower\_dBm\_reference\_us}(m) = 10 \times \log_{10}\left(\sum_{i \in \text{MEDLEYus} \cap \text{US}(m)} \left(10^{\frac{\text{PSD\_UO2}(i \times \Delta f)}{10}} \times g_i^2\right)\right)$$

where PSD_UO2($i \times \Delta f$) is the PSD measured at the U-O2 reference point with the VTU-R connected to the loop and frozen in the R-P-MEDLEY stage of initialization with the SOC in the R-IDLE state, and with the VTU-O replaced by an R_N=100 Ohm resistance terminating the loop.

NOTE – The feature to freeze a VTU in the MEDLEY stage of initialization exists solely to allow a test bed to be constructed for the purpose of measuring the H(*f*) reference values. It applies only to specific transceivers serving as the 'transmit transceiver' of the test environment and is not a requirement for compliance with this Recommendation.

The receiving VTU shall measure the SATN values under the same loop, noise, temperature and configuration settings that are used for measuring the SATN reference values.

For all downstream bands, *m*, that are active during loop diagnostics or initialization, the absolute error between SATN_D(*m*) and SATN_reference_ds(*m*) shall be equal to or smaller than 3 dB.

The accuracy requirements for downstream SATN_D(*m*) shall apply to its measurement after Initialization during showtime.

For all upstream bands, *m*, that are active during loop diagnostics or initialization, the absolute error between SATN_U(*m*) and SATN_reference_us(*m*) shall be equal to or smaller than 3 dB.

The accuracy requirements for upstream SATN_U(*m*) shall apply to its measurement after Initialization during showtime.

11.4.1.2.6 Accuracy of signal-to-noise ratio margin (*SNRM*)

For further study.

11.4.1.2.7 Accuracy of attainable net data rate (*ATTNDR*)

For further study.

11.4.1.2.8 Accuracy of actual aggregate transmit power (*ACTATP*)

The VTU-O near-end ACTATP reference value shall be defined as follows:

ACTATP_reference_UO2 = sum_over_all_frequencies [PSDps_UO2(*i*)],

where PSDps_UO2(*i*) is the downstream PSD measured at the U-O2 reference point, after initialization of the line up to the SHOWTIME state, in which state the VTU-O is frozen and the VTU-O subsequently connected to an R_N=100 Ohms.

The VTU-R near-end ACTATP reference value shall be defined as follows:

ACTATP_reference_UR2 = sum_over_all_frequencies [PSDps_UR2(*i*)]

where PSDps_UR2(*i*) is the upstream PSD measured at the U-R2 reference point, after initialization of the line up to the SHOWTIME state, in which state the VTU-R is frozen and the VTU-R subsequently connected to an R_N=100 Ohms.

NOTE 1 – The ACTATP should be measured first. Subsequently, the VTU should be frozen in SHOWTIME and the PSDps_Ux should then be measured without re-initialization.

NOTE 2 – The measurement of the PSDps_Ux involves freezing in SHOWTIME of the transceiver under test. Specification of special test modes for the transceiver under test is outside the scope of this Recommendation.

The absolute error between the VTU-O near-end ACTATP_ds and the ACTATP_reference_UO2 shall be equal to or smaller than 1.0 dB.

The sample variance of the VTU-O near-end ACTATP_ds measurements (all samples taken over a 10 minutes time interval, without line re-initialization and bit/gain-swaps in this time interval, and under the same loop, noise, temperature, and configuration settings) shall be equal to or smaller than 0.5 dB.

NOTE 3 – The ACTATP_ds samples are to be taken after sufficient time is allowed after initialization for bit and gain swaps to stabilize.

The absolute error between the VTU-R near-end ACTATP_us and the ACTATP_reference_UR2 shall be equal to or smaller than 1.0 dB.

The sample variance of the VTU-R near-end ACTATP_us measurements (all samples taken over a 10 minute time interval, without line re-initialization and bit/gain-swaps in this time interval, and under the same loop, noise, temperature, and configuration settings) shall be equal to or smaller than 0.5 dB.

NOTE 4 – The ACTATP_us samples are to be taken after sufficient time is allowed after initialization for bit and gain swaps to stabilize.

11.4.2 Configuration parameters

11.4.2.1 Transmitter-referred virtual noise PSD

This clause describes the transmitter-referred virtual noise PSD parameter TXREFVN, used only in the optional SNR margin mode SNRM_MODE = 2, SNRM_MODE = 4, and SNRM_MODE=5.

11.4.2.1.1 Definition of parameter TXREFVN

Configuration parameter TXREFVN defines the transmitter-referred virtual noise PSD to be used in determining the SNR margin.

For SNRM_MODE = 2 and SNRM_MODE = 4, the CO-MIB shall provide a TXREFVN parameter set for each utilized band.

NOTE – For SNRM_MODE=5, the TXREFVN parameter provided by the CO-MIB may be taken into account by the VTU-O to determine the initial value of TXREFSAVN (see clause 11.4.1.1.6.1.5).

The transmitter-referred virtual noise PSD in the CO-MIB shall be specified by a set of breakpoints.

Each breakpoint shall consist of a subcarrier index t_n and a noise PSD (expressed in dBm/Hz). The TXREFVN parameter for each utilized band shall be a set of breakpoints that are represented by [(t_1, PSD_1), (t_2, PSD_2), ..., (t_n, PSD_n), (t_{NBP}, PSD_{NBP})], where t_1 and t_{NBP} are, respectively, the lower and higher band edge frequencies of the band.

The subcarrier indices t_i shall be coded in the CO-MIB as unsigned integers in the range from $t_1 = roundup(f_x/Df)$ to $t_{NBP} = rounddown(f_{x+1}/Df)$, where f_x, f_{x+1} are the low and the high band separating frequencies determined by the applied band plan and specified in clause 7.1, and $Df = 4.3125$ kHz. The breakpoints shall be defined so that $t_n < t_{n+1}$ for $n = 1$ to $N - 1$; the frequency f_n corresponding to the index t_n can be found as: $f_n = t_n \times Df$. The value of Df is independent of the subcarrier spacing Δf used for DMT modulation. When the VTU operates with 8.625 kHz subcarrier spacing, all odd values of t_i shall be converted by the VTU, by rounding down to the next lower even value, and values t_1 and t_{NBP} shall be rounded (up and down, respectively) to even values.

The values for the transmitter-referred virtual noise PSD shall be coded as 8-bit unsigned integers representing virtual noise PSDs from −40 dBm/Hz (coded as 0) to −140 dBm/Hz (coded as 200), in steps of 0.5 dBm/Hz. Values from 201 to 255, inclusive, correspond to a virtual noise PSD of zero Watt/Hz (minus infinity dBm/Hz).

The maximum number of breakpoints is 32 in the downstream and 16 in the upstream.

The parameter in the downstream direction is TXREFVNds, and the parameter in the upstream direction is TXREFVNus.

NOTE – TXREFVN is configured via the ITU-T G.997.1 parameter VN, whose interpretation depends on the value of SNRM_MODE.

11.4.2.1.2 Use of parameter TXREFVN

The transmitter-referred virtual noise PSD to be used by the transceiver for calculation of the SNR margin for each subcarrier *i*, shall be obtained by linear interpolation in dB on a linear frequency scale as follows:

$$TX\_referred\_Virtual\_Noise\_PSD(i) \;=\; PSD_n \;+\; \left(PSD_{n+1} - PSD_n\right) \times \frac{\left(\frac{i^*\Delta f}{Df}\right) - t_n}{t_{n+1} - t_n} \qquad t_n < \left(\frac{i^*\Delta f}{Df}\right) \le t_{n+1}$$

where Δf is the actual subcarrier spacing used by the DMT modulation.

In downstream, the breakpoints (t_i, PSD_i) are those communicated to the VTU-R in O-SIGNATURE and are a combination of the configuration parameter TXREFVNds and the configuration parameter TXRFVNSFds as described in clause 11.4.2.5.

In upstream, the breakpoints (t_i, PSD_i) are equal to the breakpoints of the configuration parameter TXREFVNus as in the CO-MIB.

The near-end transceiver should apply the Received_Virtual_Noise_PSD (see clause 11.4.1.1.6.1.2) at the constellation decoder point (i.e., the transceiver does not need to account for DFT leakage effects from one subcarrier to another subcarrier). All effects are to be taken into account in the setting of the TXREFVN in the CO-MIB.

NOTE 1 – The above method is equivalent to the near-end transceiver calculating its bit loading using the following Virtual_Noise_SNR for the subcarrier with index *i*, at the constellation decoder (all terms are expressed in dB):

$$\text{Virtual\_Noise\_SNR}(i) = S\_tx(i) - N\_tx(i) + 20 \times \log_{10}(g_i)$$

where:

$$S\_tx(\mathrm{i}) \;= \mathrm{MREFPSD}(i)$$
$$N\_tx(i) \;= \mathrm{TX\_referred\_Virtual\_Noise\_PSD}(i)$$

and MREFPSD(*i*) is the MEDLEY reference PSD value at the far-end transmitter for the subcarrier with index *i*, obtained by interpolation of the breakpoints of the MEDLEY reference PSD (MREFPSD) information exchanged in the O-PRM and R-PRM messages during initialization.

TX_referred_Virtual_Noise_PSD(*i*) is the transmitter-referred virtual noise PSD value for subcarrier with index *i*, obtained by interpolation of the breakpoints of TXREFVN sent in the O-SIGNATURE message during initialization.

g_i is the gain adjuster for the subcarrier with index *i* as defined in clause 10.3.4.

NOTE 2 – Improper setting of TXREFVN can interact with the setting of one or more of the following parameters: maximum net data rate, downstream maximum SNR margin, impulse noise protection, and maximum interleaving delay. This interaction can result in high levels of transmit power that can lead to high crosstalk experienced by DSLs on other pairs in the same binder.

11.4.2.2 INM procedure and control parameters

This clause describes the INM procedure and associated INM control parameters.

11.4.2.2.1 INM procedure

Figure 11-5 shows the INM functional block diagram.

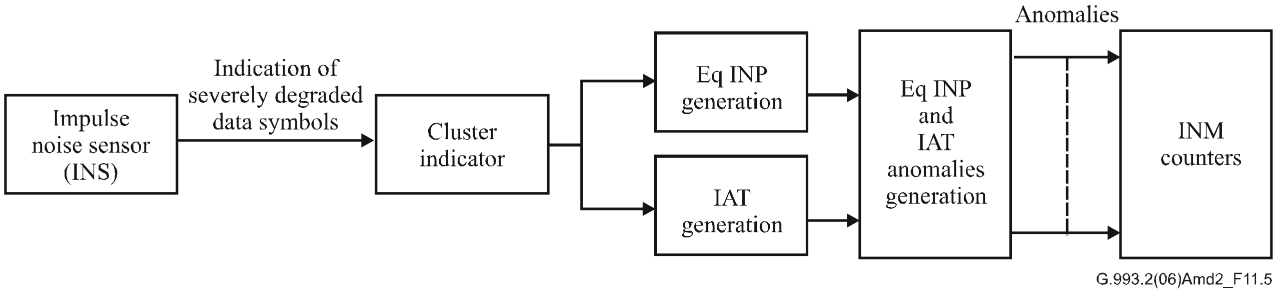

Figure 11-5 – Impulse noise monitor functional block diagram

The INS indicates whether a data symbol is severely degraded or not. A data symbol is considered to be severely degraded when it would lead to severe errors on the gamma interface when there would be no impulse noise protection (i.e., RS only used for coding gain). The implementation details for this sensor are vendor-discretionary.

NOTE 1 – Performance requirements for the INS are for further study.

If a sync symbol occurs between two data symbols (severely degraded or not), the INS shall disregard it.

NOTE 2 – If a sync symbol occurs between two groups of respectively N1 and N2 consecutive severely degraded data symbols, the two groups will form a single group of consecutive severely degraded data symbols of length N1 + N2 data symbols.

The cluster indicator indicates short groups of severely degraded data symbols as clusters. The cluster can contain a single severely degraded data symbol, a group of consecutive severely degraded data symbols, or several groups of one or more consecutive severely degraded data symbols with gaps between the groups.

The cluster indicator shall use the following rule to identify the cluster. A gap is defined as a group of non-severely degraded data symbols in-between two severely degraded data symbols. A cluster is defined as the largest group of consecutive data symbols, starting and ending with a severely degraded data symbol, containing severely degraded data symbols, separated by gaps smaller than or equal to INMCC (the cluster continuation parameter, see clause 11.4.2.2.4).

As a consequence of the above definition of a cluster, each cluster starts with a severely degraded data symbol preceded by a gap larger than INMCC and ends with a severely degraded data symbol followed by a gap larger than INMCC, while gaps inside the cluster are all smaller than or equal to INMCC.

In the Eq INP generation block, the "equivalent INP" of the cluster is generated. For each cluster, the following characteristics shall be determined:

- The impulse noise cluster length (INCL), defined as the number of data symbols from the first to the last severely degraded data symbol in the cluster;
- The impulse noise cluster degraded data symbols (INCD), defined as the number of severely degraded data symbols in the cluster;
- The impulse noise cluster number of gaps (INCG), defined as the number of gaps in the cluster, with gap as defined above.

Depending on the value of the control parameter INM_INPEQ_MODE, the equivalent INP is generated as:

- INM_INPEQ_MODE = 0: $INP\_eq = INCL$ with INMCC = 0
- INM_INPEQ_MODE = 1: $INP\_eq = INCL$ with INMCC as configured (see clause 11.4.2.2.4)

- INM_INPEQ_MODE = 2: $INP\_eq = INCD$ with INMCC as configured (see clause 11.4.2.2.4)
- INM_INPEQ_MODE = 3:

  For INCG < (8×erasuregain): $INP\_eq = \min\left(INCL, ceil\left[INCD \times \left(\frac{1}{1-\frac{INCG}{8\times erasuregain}}\right)\right]\right)$

  For INCG ≥ (8×erasuregain): $INP\_eq = INCL$

  with INMCC as configured (see clause 11.4.2.2.4)

  where the erasure gain is defined as:

  $erasuregain = \frac{INP}{INP\_no\_erasure}$ with INP, and *INP_no_erasure* as defined in clause 9.6.

  NOTE 3 – In case the bit "INP_no_erasure_required" (bit 8 in the "impulse noise protection and dynamic interleaver reconfiguration" field in Table 12-51, clause 12.3.5.2.1.1) is set, the erasure gain is equal to 1.

- INM_INPEQ_MODE = 4: In this mode, the value of INP_eq shall correspond with the VTU's own estimate (i.e., VTU-R's estimate in the downstream, VTU-O's estimate in the upstream) of the INP_min setting required to provide error-free operation for the cluster, with INMCC as configured (see clause 11.4.2.2.4). The method of computation of the VTU's own estimate is vendor-discretionary. For INM_INPEQ_MODE = 4 only, if INMCC is set to 64, the VTU shall use its own method for cluster indication. If INMCC < 64, the VTU shall use the cluster indicator as described in this clause for the INM_INPEQ_MODE = 1, 2 and 3.

Anomalies are generated for several values of INP_eq, as defined in clause 11.3.4.1. The counters of these anomalies represent the INP_eq histogram.

In the IAT generation block, the inter-arrival time (IAT) is generated as the number of data symbols from the start of a cluster to the start of the next cluster. If sync symbols occur between two clusters, they shall not be counted in the IAT. Anomalies are generated for several ranges of inter arrival time, as defined in clause 11.3.4.3. The counters of these anomalies represent the IAT histogram.

For every data symbol, the total measurement count INMAME is increased by 1.

11.4.2.2.2 Definition of parameter INMIATO

Configuration parameter INMIATO defines the INM inter-arrival time offset for the IAT anomaly generation in order to determine in which bin of the inter-arrival time histogram the IAT is reported (see clause 11.3.4.3).

The CO MIB shall provide the value for the INMIATO parameter. The parameter in the downstream direction is INMIATOds, and the parameter in the upstream direction is INMIATOus.

The valid values for INMIATO in both directions range from 3 to 511 DMT symbols in steps of 1 DMT symbol. If the VTU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the VTU-R shall use a default value of INMIATOds = 3. During showtime, this value may be overwritten by the VTU-O using an INM facility command defined in clause 11.2.3.13. A link state transition shall not affect the INMIATOds value (e.g., not reset the value to the default value).

The VTU-O shall use the current value of INMIATOus stored in the CO MIB.

11.4.2.2.3 Definition of parameter INMIATS

Configuration parameter INMIATS defines the INM inter-arrival time step for the IAT anomaly generation in order to determine in which bin of the inter-arrival time histogram the IAT is reported (see clause 11.3.4.3).

The CO MIB shall provide the value for the INMIATS parameter. The parameter in the downstream direction is INMIATSds, and the parameter in the upstream direction is INMIATSus.

The valid values for INMIATS range from 0 to 7 in steps of 1. If the VTU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the VTU-R shall use a default value of INMIATSds = 0. During showtime, this value may be overwritten by the VTU-O using an INM facility command defined in clause 11.2.3.13. A link state transition shall not affect the INMIATSds value (e.g., not reset the value to the default value).

The VTU-O shall use the current value of INMIATSus stored in the CO MIB.

11.4.2.2.4 Definition of parameter INMCC

Configuration parameter INMCC defines the INM cluster continuation value to be used in the cluster indication process described in clause 11.4.2.2.1. If INM_INPEQ_MODE = 0, INMCC is equal to zero, independent of the CO MIB setting. If INM_INPEQ_MODE > 0, the CO MIB shall provide the value for the INMCC parameter. The parameter in the downstream direction is INMCCds, and the parameter in the upstream direction is INMCCus.

The valid values for INMCC range from 0 to 64 DMT symbols in steps of 1 DMT symbol. If the VTU supports the INM facility, it shall support INMCC = 0. If the VTU supports the INM facility, and supports any INM_INPEQ_MODE > 0, it shall support all valid values.

Upon entering the first showtime after power-up, the VTU-R shall use a default value of INMCCds = 0. During showtime, this value may be overwritten by the VTU-O using an INM facility command defined in clause 11.2.3.13.

A link state transition shall not affect the INMCCds value (e.g., not reset the value to the default value).

The VTU-O shall use the current value of INMCCus stored in the CO MIB.

11.4.2.2.5 Definition of parameter INM_INPEQ_MODE

Configuration parameter INM_INPEQ_MODE defines the way of computation of equivalent INP, as defined in clause 11.4.2.2.1. The CO MIB shall provide the value for the INM_INPEQ_MODE parameter. The parameter in the downstream direction is INM_INPEQ_MODEds, and the parameter in the upstream direction is INM_INPEQ_MODEus.

The valid values for INM_INPEQ_MODE are 0, 1, 2, 3 and 4. If the VTU supports the INM facility, it shall support INM_INPEQ_MODE = 0. All other modes are optional. If the VTU supports any INM_INPEQ_MODE > 0, it shall support at least INM_INPEQ_MODE = 1, 2 and 3.

Upon entering the first showtime after power-up, the VTU-R shall use a default value of INM_INPEQ_MODEds = 0. During showtime, this value may be overwritten by the VTU-O using an INM facility command defined in clause 11.2.3.13.

A link state transition shall not affect the INM_INPEQ_MODE value (e.g., not reset the value to the default value).

The VTU-O shall use the current value of INM_INPEQ_MODEus stored in the CO MIB.

11.4.2.2.6 Definition of parameter INM_INPEQ_FORMAT

Configuration parameter INM_INPEQ_FORMAT defines the way the scale is configured for the INM_INPEQ histogram, as defined in clause 11.3.4.1. The CO MIB shall provide the value for the INM_INPEQ_FORMAT parameter. The parameter in the downstream direction is INM_INPEQ_FORMATds, and the parameter in the upstream direction is INM_INPEQ_FORMATus.

The valid values for INM_INPEQ_FORMAT are 0 and 1. If the VTU supports the INM facility, it shall support INM_INPEQ_FORMAT = 0 and may support INM_INPEQ_FORMAT = 1.

Upon entering the first showtime after power-up, the VTU-R shall use a default value of INM_INPEQ_FORMAT = 0. During showtime, this value may be overwritten by the VTU-O using an INM facility command defined in clause 11.2.3.13 if the VTU-R supports INM_INPEQ_FORMAT = 1.

A link state transition shall not affect the INM_INPEQ_FORMAT value (e.g., not reset the value to the default value).

The VTU-O shall use the current value of INM_INPEQ_FORMATus stored in the CO-MIB.

11.4.2.3 Receiver-referred virtual noise PSD

This clause describes the receiver-referred virtual noise PSD parameter RXREFVN, used only in the optional SNR margin mode SNRM_MODE = 3 and SNRM_MODE=4.

11.4.2.3.1 Definition of parameter RXREFVN

Configuration parameter RXREFVN defines the receiver-referred virtual noise PSD to be used in determining the SNR margin.

For SNRM_MODE=3 and SNRM_MODE=4, the CO-MIB shall provide a RXREFVN parameter set for each utilized band.

The receiver-referred virtual noise PSD in the CO-MIB shall be specified by a set of breakpoints. Each breakpoint shall consist of a subcarrier index t_n and a noise PSD value (expressed in *dBm/Hz*). The RXREFVN parameter for each band shall be a set of breakpoints that are represented by $[(t_1, PSD_1), (t_2, PSD_2), \dots, (t_n, PSD_n), \dots, (t_{NBP}, PSD_{NBP})]$, where t_1 and t_{NBP} are, respectively, the lower and higher band edge frequencies of the band. The VTU shall ignore any frequency information that does not belong to the utilized (upstream) bands.

The subcarrier indices t_i shall be coded in the CO-MIB as unsigned integers in the range from $t_1 = roundup(f_x/Df)$ to $t_{NBP} = rounddown(f_{x+1}/Df)$, where f_x, f_{x+1} are the low and the high band separating frequencies determined by the applied band plan and specified in clause 7.1, and $Df = 4.3125$ kHz. The breakpoints shall be defined so that $t_n < t_{n+1}$ for $n = 1$ to $N - 1$; the frequency f_n corresponding to the index t_n can be found as: $f_n = t_n \times Df$. The value of Df is independent of the subcarrier spacing Δf used for DMT modulation. When the VTU operates with 8.625 kHz subcarrier spacing, all odd values of t_i shall be converted by the VTU, by rounding down to the next lower even value, and values t_1 and t_{NBP} shall be rounded (up and down, respectively) to even values.

The values for the virtual noise PSD shall be coded as 8-bit unsigned integers representing virtual noise PSD values from −40 dBm/Hz (coded as 0) to −140 dBm/Hz (coded as 200), in steps of 0.5 dBm/Hz. Values from 201 to 255, inclusive, correspond to a virtual noise PSD of zero Watt/Hz (minus infinity dBm/Hz).

The maximum number of breakpoints is 16.

The parameter in the upstream direction is RXREFVNus.

NOTE – RXREFVN is configured via the ITU-T G.997.1 parameter VN, whose interpretation depends on the value of SNRM_MODE.

11.4.2.3.2 Use of parameter RXREFVN

For each frequency band, the receiver-referred virtual noise PSD, for each subcarrier *i*, shall be obtained by linear interpolation in *dB* on a linear frequency scale as follows:

$$RX\_referred\_Virtual\_Noise\_PSD(i) = PSD_n + \left(PSD_{n+1} - PSD_n\right) \times \frac{\left(\frac{i*\Delta f}{Df}\right) - t_n}{t_{n+1} - t_n} \qquad t_n < \left(\frac{i*\Delta f}{Df}\right) \le t_{n+1}$$

where Δf is the actual subcarrier spacing used by the DMT modulation.

In SNRM_MODE=3, the breakpoints (t_i, PSD_i) are equal to the breakpoints of the configuration parameter RXREFVNus as in the CO-MIB.

In SNRM_MODE=4, the breakpoints (t_i, PSD_i) are a combination of the configuration parameter RXREFVNus and the configuration parameter RXRFVNSFus as described in clause 11.4.2.4.

The near-end transceiver should apply the Received_Virtual_Noise_PSD (see clause 11.4.1.1.6.1.3) over the upstream frequencies at the constellation decoder point (i.e., the transceiver does not need to account for DFT leakage effects from one subcarrier to another subcarrier). All effects are to be taken into account in the setting of the RXREFVN in the CO-MIB.

NOTE 1 – The above method is equivalent to the near-end transceiver calculating its bit loading using the following Virtual_Noise_SNR for the subcarrier with index i, at the constellation decoder (all terms are expressed in dB):

Virtual_Noise_SNR(i) = $S\_rx(i) – N\_rx(i)$

where:

S_tx(i) = Actual_Received_Signal_PSD = $|H_{RXfilter}(f)|^2$ + Actual_Received_Signal_at_U_interface

$N\_tx(i)$ = Received_Virtual_Noise_PSD = $|H_{RXfilter}(f)|^2$ + RXREFVN + RXREFVNSF

NOTE 2 – Improper setting of RXREFVN and RXREFVNSF can interact with the setting of one or more of the following parameters: maximum net data rate, downstream maximum SNR margin, impulse noise protection, and maximum interleaving delay. This interaction can result in high levels of transmit power that can lead to high crosstalk experienced by DSLs on other pairs in the same binder.

11.4.2.4 Receiver-referred virtual noise scaling factor

This clause describes the receiver-referred virtual noise scaling factor parameter RXREFVNSF, used only in the optional SNR margin mode SNRM_MODE = 4.

Configuration parameter RXREFVNSF defines the receiver-referred virtual noise scaling factor to be used together with the receiver-referred virtual noise PSD in determining the SNR margin.

The CO-MIB shall provide an upstream RXREFVNSF parameter when SNRM_MODE = 4.

The values for the receiver-referred virtual noise PSD scaling factor shall be coded as 8-bit signed integers representing scaling factors from −64.0 dB (coded as −128) to 63.5 dB (coded as 127), in steps of 0.5 dB.

The parameter in the upstream direction is RXREFVNSFus.

The VTU-O shall combine the value of the configuration parameter RXREFVNSFus as in the CO-MIB with the value of the configuration parameter RXREFVNus as in the CO-MIB to a control parameter RXREFVNus as follows:

Control parameter RXREFVNus in *dBm*/*Hz* = configuration parameter RXREFVNus in *dBm*/*Hz* + configuration parameter RXREFVNSFus in *dB*.

11.4.2.5 Transmitter-referred virtual noise scaling factor

This clause describes the transmitter-referred virtual noise scaling factor parameter TXREFVNSF, used only in the optional SNR margin mode SNRM_MODE = 4.

Configuration parameter TXREFVNSF defines the transmitter-referred virtual noise scaling factor to be used together with the transmitter-referred virtual noise PSD in determining the SNR margin.

The CO-MIB shall provide a TXREFVNSF parameter for downstream when SNRM_MODE = 4.

The values for the receiver-referred virtual noise PSD scaling factor shall be coded as 8-bit signed integers representing scaling factors from −64.0 dB (coded as –128) to +63.5 dB (coded as +127), in steps of 0.5 dB.

The parameter in the downstream direction is TXREFVNSFds.

The VTU-O shall combine the value of the configuration parameter TXREFVNSFds as in the CO-MIB with the value of the configuration parameter TXREFVNds as in the CO-MIB to a control parameter TXREFVNds as communicated to the VTU-R in O-SIGNATURE as follows:

Control parameter TXREFVNds in *dBm*/*Hz* = *min*(*max*((configuration parameter TXREFVNds in *dBm*/*Hz* + configuration parameter TXREFVNSFds in *dB*), –140 *dBm*/*Hz*), –40 *dBm*/*Hz*).

11.4.2.6 Re-Initialization policy parameters

11.4.2.6.1 Re-Initialization policy selection (RIPOLICY)

The RIPOLICY configuration parameter (see clause 7.3.1.1.12 of [ITU-T G.997.1]) indicates which policy shall be applied to determine the triggers for re-initialization in a specific direction. The parameter in the downstream direction is RIPOLICYds, and the parameter in the upstream direction is RIPOLICYus.

The control parameter $RIpolicyds_n$ shall have the same value as the configuration parameter RIPOLICYds in the CO-MIB for all downstream bearer channels. However, if configuration parameter RIPOLICYds=1 and the VTU-R indicates during ITU-T G.993.2 initialization in R-MSG 2 (see clause 12.3.5.2.2.1) that $RIpolicyds_n$=1 is not supported for a particular bearer channel, the VTU-O shall fall back to indicate $RIpolicyds_n$=0 in O-TPS (see clause 12.3.5.2.1.2) for that bearer channel.

The control parameter $RIpolicyus_n$ shall have the same value as the configuration parameter RIPOLICYus in the CO-MIB for all upstream bearer channels. However, if $RIpolicyus_n$=1 is not supported by the VTU-O for a particular bearer channel, the VTU-O shall fall back to $RIpolicyus_n$=0 for that bearer channel.

The valid values for RIPOLICY are 0 and 1.

11.4.2.6.2 REINIT_TIME_THRESHOLD

Configuration parameter REINIT_TIME_THRESHOLD (see clause 7.3.1.1.13 of [ITU-T G.997.1]) defines the threshold for re-initialization based on SES, to be used by the VTU receiver when Re-Initialization Policy 1 is used (see clause 12.1.4).

The downstream and upstream values of REINIT_TIME_THRESHOLD shall be configured in the CO-MIB. The parameter in the downstream direction is REINIT_TIME_THRESHOLDds, and the parameter in the upstream direction is REINIT_TIME_THRESHOLDus.

The control parameter *REINIT_TIME_THRESHOLDds* conveyed in O-MSG 1 shall have the same value as the configuration parameter REINIT_TIME_THRESHOLDds in the CO-MIB. The control parameter *REINIT_TIME_THRESHOLDus* shall have the same value as the configuration parameter REINIT_TIME_THRESHOLDus in the CO-MIB.

The value shall be coded as an unsigned integer representing the maximum number of SES. The valid range is from 5 to 31.

11.4.2.7 MAXDELAYOCTET-split (MDOSPLIT)

The line configuration parameter MAXDELAYOCTET-split (MDOSPLIT) defines the percentage of the MAXDELAYOCTET allocated to the downstream direction. All of the remaining MAXDELAYOCTET shall be allocated for use in the upstream direction.

MAXDELAYOCTET is the parameter "aggregate interleaver and de-interleaver delay", in octets, specified in Table 6-1 for the profile. A single MAXDELAYOCTET-split parameter is defined. The allocated number of octets shall be the total allocated to a given transmission direction for interleaving, and for all latency paths.

Configuration of a split in percentage would result in fractional octets. Therefore, the following rounding rule shall apply ($\lceil x \rceil$ denotes rounding to the higher integer):

$$\text{MAXDELAYOCTET\_DS} = \lceil \text{MDOSPLIT} \times \text{MAXDELAYOCTET} \rceil$$

$$\text{MAXDELAYOCTET\_US} = \text{MAXDELAYOCTET} - \text{MAXDELAYOCTET\_DS}$$

The number of octets used in downstream direction (aggregated over both transceivers) shall be no higher than MAXDELAYOCTETS_DS. The number of octets used in upstream direction (aggregated over both transceivers) shall be no higher than MAXDELAYOCTETS_US.

The sum of the max_delay_octet values specified in O-PMS (see clause 12.3.5.2.1.3) shall be limited to:

$$\text{max\_delay\_octet}_{\text{DS},0} + \text{max\_delay\_octet}_{\text{DS},1} \le \text{MAXDELAYOCTETS\_DS}$$

$$\text{max\_delay\_octet}_{\text{US},0} + \text{max\_delay\_octet}_{\text{US},1} \le \text{MAXDELAYOCTETS\_US}$$

The above constraints depend only on the MAXDELAYOCTET-split parameter configured through the CO-MIB. The VTU-O shall not re-distribute the allocated max_delay_octets over downstream and upstream directions based on line conditions. As a result, some octets allocated to a given direction may remain unused.

In case of dual latency (i.e., two bearer channels each in a different latency path), the VTU-O shall determine the max_delay_octet values allocated to each of the latency paths to optimize the data rate within the constraint given by the Rate Adaptation Ratio parameter defined in clause 7.3.2.1.4 of [ITU-T G.997.1].

MDOSPLIT shall be expressed as a percentage, with valid range from 5 percent to 95 percent inclusive, in steps of 1 percent. The value 0% is valid only if the maximum downstream interleaving delay is configured with the special value S1 (see clause 7.3.2.2 of [ITU-T G.997.1]). The value 100% is valid only if the maximum upstream interleaving delay is configured with the special value S1 (see clause 7.3.2.2 of [ITU-T G.997.1]). A special value shall indicate that the VTU-O is allowed to use a vendor discretionary algorithm to determine the max_delay_octet values specified in O-PMS.

NOTE – A special value is introduced to ensure backward compatibility.

11.4.2.8 ATTNDR_MAXDELAYOCTET-split (ATTNDR_MDOSPLIT)

The line configuration parameter ATTNDR_MAXDELAYOCTET-split (ATTNDR_MDOSPLIT) defines the percentage of the MAXDELAYOCTET allocated to the downstream direction to be used in the calculation of the ATTNDR. All of the remaining MAXDELAYOCTET shall be allocated for use in the upstream direction.

MAXDELAYOCTET is the parameter "aggregate interleaver and de-interleaver delay", in octets, specified in Table 6-1 for the profile. A single ATTNDR_MDOSPLIT parameter is defined. The allocated number of octets shall be the total allocated to a given transmission direction for interleaving, and for all latency paths.

Configuration of a split in percentage would result in fractional octets. Therefore, the following rounding rule shall apply ($\lceil x \rceil$ denotes rounding to the higher integer):

$$ATTNDR\_MAXDELAYOCTET\_DS = \lceil ATTNDR\_MDOSPLIT \times MAXDELAYOCTET \rceil$$

$$ATTNDR\_MAXDELAYOCTET\_US = MAXDELAYOCTET - ATTNDR\_MAXDELAYOCTET\_DS$$

The number of octets used in downstream direction (aggregated over both transceivers) shall be no higher than ATTNDR_MAXDELAYOCTET_DS. The number of octets used in upstream direction (aggregated over both transceivers) shall be no higher than ATTNDR_MAXDELAYOCTET_US.

If the ATTNDR calculation uses single latency mode, the ATTNDR_max_delay_octet values shall be defined as:

$$ATTNDR\_max\_delay\_octet_{DS,0} = ATTNDR\_MAXDELAYOCTET\_DS$$

$$ATTNDR\_max\_delay\_octet_{US,0} = ATTNDR\_MAXDELAYOCTET\_US$$

where the $ATTNDR\_max\_delay\_octet_{DS,0}$ specifies the maximum of $delay\_octet_{DS,0}$ that the VTU-R shall assume in the calculation of ATTNDR in downstream (see clause 12.3.5.2.1.3), and the $ATTNDR\_max\_delay\_octet_{US,0}$ specifies the maximum of $delay\_octet_{US,0}$ that the VTU-O shall assume in the calculation of ATTNDR in upstream.

If the ATTNDR calculation uses single latency with ROC mode, the ATTNDR_max_delay_octet values shall be defined as:

$$ATTNDR\_max\_delay\_octet_{DS,1} = ATTNDR\_MAXDELAYOCTET\_DS - delay\_octet_{DS,0}$$

$$ATTNDR\_max\_delay\_octet_{US,1} = ATTNDR\_MAXDELAYOCTET\_US - delay\_octet_{US,0}$$

where the $delay\_octet_{DS,0}$ and $delay\_octet_{US,0}$ values correspond to the actual configuration of the latency path #0 as applicable at the instant of ATTNDR calculation,

and the $ATTNDR\_max\_delay\_octet_{DS,1}$ specifies the maximum of $delay\_octet_{DS,1}$ that the VTU-R shall assume in the calculation of ATTNDR in downstream (see clause 12.3.5.2.1.3),

and the $ATTNDR\_max\_delay\_octet_{US,1}$ specifies the maximum of $delay\_octet_{US,1}$ that the VTU-O shall assume in the calculation of ATTNDR in upstream.

The above constraints depend only on the ATTNDR_MDOSPLIT parameter configured through the CO-MIB.

ATTNDR_MDOSPLIT shall be expressed as a percentage, with valid range from 5 per cent to 95 per cent inclusive, in steps of 1 per cent. The value 0% is valid only if the maximum downstream interleaving delay is configured with the special value S1 (see clause 7.3.2.2 of [ITU-T G.997.1]). The value 100% is valid only if the maximum upstream interleaving delay is configured with the special value S1 (see clause 7.3.2.2 of [ITU-T G.997.1]). A special value shall indicate that the VTU-O is allowed to use a vendor discretionary algorithm to determine the ATTNDR_max_delay_octet values specified in O-PMS.

11.5 Data gathering function

The data gathering function stores records in a buffer, with each record entered in the buffer upon occurrence of a particular event type. These records may be used for diagnosing troubles and analysing functions. Different event types are defined; these are relevant to the operation of the line and are useful for determining what is functioning properly or what is not. Each event is recorded with a timestamp, the ID of the event type, and additional event data as defined for each event type.

The VTU-O may optionally support the data gathering function. If supported, then the data gathering function shall comply with this clause, as applicable to the VTU-O.

The VTU-R may optionally support the data gathering function. If supported, then the data gathering function shall comply with this clause, as applicable to the VTU-R.

The data gathering function uses both the ITU-T G.993.2 transceiver functionality and the ITU-T G.997.1 functionality. Support of some ITU-T G.997.1 functionality is required to implement this clause. The data gathering function shall include a buffer of *logging_depth* records. Each record shall include the following three fields: Event Timestamp, Event Type and Event Data.

For each event type, a logging depth is defined as the minimum number of records of the event type that shall be kept in the buffer when deleting entries of this event type in case of buffer full condition. Per event type, the logging depth shall be configurable as a parameter derived from the configuration parameter $logging\_depth\_event\_percentage_i$ as follows:

$$logging\_depth\_event_i = \lfloor (logging\_depth\_event\_percentage_i / 100\%) \times logging\_depth \rfloor$$

The sum over all event types of the configured $logging\_depth\_event\_percentage_i$ shall not exceed 100%, i.e.,

$$\sum_i logging\_depth\_event\_percentage_i \leq 100\ \%$$

NOTE – If the sum of $logging\_depth\_event_i$ is smaller than *logging_depth*, then the remaining part of the buffer is used as a pool that can be used for any event type, within the buffer management rules defined in clause 11.5.2.

The control parameter, $logging\_depth\_event_i$, is the minimum logging depth for event type *i* in number of 6-byte records, i.e., $logging\_depth\_event_i$ is the minimum number of records of event type *i* that shall be kept in the event buffer of the VTU. This is logically the same as having a separate buffer for each event type, with each of these buffers able to store at least $logging\_depth\_event_i$ records.

The maximum depth of the entire event buffer is reported by the VTU as *logging_depth*.

The parameter *logging_depth* is defined separately for the VTU-R and the VTU-O respectively as *logging_depth_R* and *logging_depth_O*. The reporting parameter *logging_depth_R is* the maximum depth of the buffer stored at the VTU-R and is reported in the CO-MIB as LOGGING_DEPTH_R (see clause 7.5.3.2 of [ITU-T G.997.1]). The reporting parameter *logging_depth_O is* the maximum depth of the buffer stored at the VTU-O and is reported in the CO-MIB as LOGGING_DEPTH_O (see clause 7.5.3.1 of [ITU-T G.997.1]).

The control parameter $logging\_depth\_event\_percentage_i$ is defined separately for the VTU-R and the VTU-O respectively as $logging\_depth\_event\_percentage\_R_i$ and $logging\_depth\_event\_percentage\_O_i$. The control parameter $logging\_depth\_event\_percentage\_R_i$ shall have a value equal to the CO-MIB configuration parameter $LOGGING\_DEPTH\_EVENT\_PERCENTAGE\_R_i$ (see clause 7.3.6.2 of [ITU-T G.997.1]). The control parameter $logging\_depth\_event\_percentage\_O_i$ shall have a value equal to the CO-MIB configuration parameter $LOGGING\_DEPTH\_EVENT\_PERCENTAGE\_O_i$ (see clause 7.3.6.1 of [ITU-T G.997.1]).

At the VTU-O and VTU-R, the record length shall be fixed at 6 bytes (consisting of 4 bytes for the timestamp, 1 byte for the event type, and 1 byte for the event data).

At the VTU-O, the timestamps as presented at the Q-interface (towards the MIB) shall be the number of seconds since Jan 1, 1900 as defined in [IETF RFC 5905] (NTP) coded in a 32-bit format. At the VTU-R, the timestamp as presented towards the U-interface (on the DSL) shall be the number of seconds since the last power cycle (mains power turned off/on) of the VTU-R, coded in a 32-bit format (with accuracy of time clock within ±50 ppm). The VTU-R timestamp shall be recorded in the CO-MIB in the same format as it presented toward the U-interface.

A single record of each event type shall be stored once each second, unless no event of this type occurred during that second. In case an event (with a particular event ID, as per Table 11-43) has occurred multiple times during a one-second period, a single event shall be recorded with the timestamp for that one second period, and within this record the number of occurrences of the event may be reflected as part of the event data.

The following subclauses identify which event types shall be supported by the VTU-O and/or VTU-R. The reporting of all supported event types is mandatory, i.e., there are no event types that are optionally supported.

Clause 11.5.2 defines the buffer management rules.

Clause 11.2.3.16 defines the mechanism for the VTU-O to retrieve over the eoc channel the data gathered in the VTU-R buffer.

The logging depth that is applied for reporting the VTU-R buffer over the eoc channel is *logging_depth_reporting*.

The VTU-O shall set *act_logging_depth_reporting_R* in the eoc configuration command message. The *act_logging_depth_reporting_R* shall be the minimum of:

- The *logging_depth_R* (indicated by the VTU-R in the eoc configuration response message);
- The maximum depth of the buffer at the VTU-O that stores the records originated at the VTU-R as reported over the eoc channel;
- The CO-MIB configuration parameter LOGGING_DEPTH_REPORTING_R (see clause 7.3.6.4 of [ITU-T G.997.1]).

When sending a first configuration command, the VTU-O is not aware of the *logging_depth_R*. If the configuration command message indicates an *act_logging_depth_reporting_R* that is higher than *logging_depth_R*, then the VTU-R shall send a NACK response, and the VTU-O shall send a new configuration command to satisfy the above three conditions.

The control parameter *act_logging_depth_reporting_R* is reported in the CO-MIB as ACT_LOGGING_DEPTH_REPORTING_R (see clause 7.5.3.4 of [ITU-T G.997.1]).

The VTU-O shall set *act_logging_depth_reporting_O*. The *act_logging_depth_reporting_O* shall be the minimum of:

- the logging_depth_O;
- the CO-MIB configuration parameter LOGGING_DEPTH_REPORTING_O (see clause 7.3.6.3 of [ITU-T G.997.1]).

The control parameter *act_logging_depth_reporting_O* is reported in the CO-MIB as ACT_LOGGING_DEPTH_REPORTING_O (see clause 7.5.3.3 of [ITU-T G.997.1]).

If act*_logging_depth_reporting_R* is strictly less than *logging_depth_R*, then the VTU-R may store more than *act_logging_depth_reporting_R* records. In this case however, the VTU-R shall virtually reduce the number of stored records to no more than *act_logging_depth_reporting_R* records by applying the buffer management rules of clause 11.5.2 before reporting records over the eoc channel. For the data gathering functionality defined in this Recommendation, the term "VTU-R buffer" shall refer to the storage of up to *act_logging_depth_reporting_R* records.

Upon a power cycle, the VTU-R:

- shall reset the entire buffer by setting all records to be dummy records, where a dummy record is defined as a record which has all fields set to the value 0; and,
- shall set *Nnack* = 0; and,
- shall start data gathering with a reset of the timestamp to value 0, *act_logging_depth_reporting_R* = *logging_depth_R* and $logging\_depth\_percentage\_R_i = 0$ for all event types (i.e., the VTU-R buffer operates as a pure FIFO until the *logging_depth_reporting* and $logging\_depth\_event\_percentage\_R_i$ are configured).

A VTU-R supporting the data gathering functionality shall support the "previous-loss-of-power" (PLPR) flag.

The records originated at the VTU-R (and reported to the VTU-O over the eoc channel) are contained in the CO-MIB parameter EVENT_TRACE_BUFFER_R (see clause 7.5.3.6 of [ITU-T G.997.1]). The maximum number of records contained in this CO-MIB parameter is ACT_LOGGING_DEPTH_REPORTING_R (see clause 7.5.3.4 of [ITU-T G.997.1]). In the case that the number of records contained in this CO-MIB parameter is less than ACT_LOGGING_DEPTH_REPORTING_R, the remainder of the EVENT_TRACE_BUFFER_R shall be set to all dummy records.

The records originated at the VTU-O are contained in the CO-MIB parameter EVENT_TRACE_BUFFER_O (see clause 7.5.3.5 of [ITU-T G.997.1]). The maximum number of records contained in this CO-MIB parameter is ACT_LOGGING_DEPTH_REPORTING_O (see clause 7.5.3.3 of [ITU-T G.997.1]). In the case that the number of records contained in this CO-MIB parameter is less than ACT_LOGGING_DEPTH_REPORTING_O, the remainder of the EVENT_TRACE_BUFFER_O shall be set to all dummy records.

11.5.1 Event types and event data

Table 11-43 lists the event types defined for the data gathering function and the identifier that shall be used in the record to identify the event type. Table 11-43 also indicates whether the event type shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support" and O indicates "optional support").

Table 11-43 – List of event types

Event type	ID	VTU-O	VTU-R	Definition
Dummy	00_{16}	M	M	See clause 11.5
End of Showtime	01_{16}	M	M	See clause 11.5.1.1
Previous End of Showtime	02_{16}	M	M	See clause 11.5.1.2
Failed Init	03_{16}	M	M	See clause 11.5.1.3
Successful Init	04_{16}	M	M	See clause 11.5.1.4
Downstream Init Net Data Rate	05_{16}	M	M	See clause 11.5.1.5
Upstream Init Net Data Rate	06_{16}	M	M	See clause 11.5.1.6
Line Failure	07_{16}	M	M	See clause 11.5.1.7
CRC-8 Anomalies	08_{16}	M	M	See clause 11.5.1.8
OLR	09_{16}	M	M	See clause 11.5.1.9
Bitswap	$0A_{16}$	M	M	See clause 11.5.1.10
Downstream Net Data Rate After Successful SRA	$0B_{16}$	M	M	See clause 11.5.1.11
Upstream Net Data Rate After Successful SRA	$0C_{16}$	M	M	See clause 11.5.1.12
Downstream Net Data Rate After Successful SOS	$0D_{16}$	M	M	See clause 11.5.1.13
Upstream Net Data Rate After Successful SOS	$0E_{16}$	M	M	See clause 11.5.1.14
Defect	$0F_{16}$	M	M	See clause 11.5.1.15
Retransmission event	10_{16}	M	M	See clause 11.5.1.16
Reserved for use by ITU-T	11_{16} to $7F_{16}$			

Table 11-43 – List of event types

Event type	ID	VTU-O	VTU-R	Definition
Vendor-discretionary event types	80_{16} to FF_{16}			
NOTE – Event types that are optional in this Recommendation are mandatory for the data gathering function only if the particular event type is supported by the VTU.				

11.5.1.1 Event type "End of Showtime"

The event type "End of Showtime" shall occur when the VTU transitions from the SHOWTIME state to the SILENT state (see VTU-O state diagram in Figure 12-2 and VTU-R state diagram in Figure 12-3). The related event data shall identify the trigger issued at VTU-O or VTU-R for this transition.

Table 11-44 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-44 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support" and RI1 indicates "mandatory only if $RIpolicy_n$=1".)

Table 11-44 – Event data for event type "End of Showtime"

Event data	ID	VTU-O	VTU-R	Definition
Reserved for future use	00_{16}			
VTU-O host triggered re-initialization	01_{16}	M	M	L3_request and L3_grant in Figure 12-2
VTU-R host triggered re-initialization	02_{16}	M	M	L3_request and L3_grant in Figure 12-3
Persistent ne-LOS	03_{16}	M	M	See clause 12.1.4
Persistent ne-LOF	04_{16}	M	M	See clause 12.1.4
Contiguous ne-SES	05_{16}	M	M	See Table 12-4 (Note 2)
Persistent ne-LOM	06_{16}	M	M	See clause 12.1.4
Persistent ne-TPS-TC out-of-sync	07_{16}	M	M	See clause 12.1.4
MAX-SOS	08_{16}	M	M	See Table 12-3
SOS rate low timeout	09_{16}	M	M	See Table 12-3
EOC no response	$0A_{16}$	M	M	See Table 12-3
Other vendor specific VTU-O or VTU-R near-end conditions declaring a high_BER event	$0B_{16}$	M	M	See Table 12-3 (Note 1)
Other vendor specific VTU-O or VTU-R far-end conditions declaring a high_BER event	$0C_{16}$	M	M	See Table 12-3 (Note 1)
High BER-hs event as defined for $RIpolicy_n$=1	$0D_{16}$	RI1	RI1	See Table 12-4
Reserved for use by ITU-T	$0E_{16}$ to FF_{16}			
NOTE 1 – This event data ID shall only be used with $RIpolicy_n$ = 0. NOTE 2 – Only if $RIpolicy_n$=1, Table 12-4 defines this event data, otherwise it is vendor specific.				

11.5.1.2 Event type "Previous End of Showtime"

The event type "Previous End of Showtime" shall occur when the VTU enters the SHOWTIME state (see VTU-O state diagram in Figure 12-2 and VTU-R state diagram in Figure 12-3) and the VTU-R has an initialization flag set to 1 (see clause 11.2.3.6). The purpose of this event is to indicate that the previous End of Showtime was due to a power cycle or host reinit at the VTU-R. This is important because the VTU-R's timestamp is reset when either of these actions occurs. The related event data shall represent the initialization flag set at the VTU-R (PLPR and PHPR as indicated by the VTU-R in the inventory commands and response, see clause 11.2.3.6).

Table 11-45 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-45 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support" and NA indicates "not applicable").

Table 11-45 – Event data for event type "Previous End of Showtime"

Event data	ID	VTU-O	VTU-R	Definition
Reserved for future use	00_{16}			
Previous-loss-of-power (PLPR)	01_{16}	NA	M	See clause 11.2.3.6
Previous-host-reinit (PHRI)	02_{16}	NA	M	See clause 11.2.3.6
Reserved for use by ITU-T	03_{16} to FF_{16}			
NOTE – The timestamp for this event type is recorded immediately after the VTU-R is re-initialized.				

11.5.1.3 Event type "Failed Init"

The event type "Failed Init" shall occur when the VTU transitions from the INIT/TRAIN or INIT/HS state to the SILENT state (see VTU-O state diagram in Figure 12-2 and VTU-R state diagram in Figure 12-3). The related event data shall identify the last initialization phase the VTU-O or VTU-R entered before the initialization failed, e.g., channel discovery phase.

Table 11-46 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-46 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support"). The octet specifying the event data for this event type is split in half, with the four MSB defining in which phase the initialization failed, and the four LSB specifying the initialization failure cause. The initialization failure cause shall be as defined in clause 7.5.1.6 of [ITU-T G.997.1] and shall be coded as a 4-bit unsigned integer (as shown in Table 11-45).

Table 11-46 – Event data for event type "Failed Init"

Event data	Four MSB of ID	VTU-O	VTU-R	Definition
Reserved for use by ITU-T	0_{16}			
ITU-T G.994.1 phase	1_{16}	M	M	See Figure 12-5
Channel discovery phase	2_{16}	M	M	See Figure 12-5
Training phase	3_{16}	M	M	See Figure 12-5
Channel analysis and exchange phase	4_{16}	M	M	See Figure 12-5
Reserved for future use	5_{16} to F_{16}			
Event data	**Four LSB of ID**	**VTU-O**	**VTU-R**	**Definition**
Successful	0_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]
Configuration error	1_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]
Configuration not feasible on the line	2_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]
Communication problem	3_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]

Table 11-46 – Event data for event type "Failed Init"

Event data	Four MSB of ID	VTU-O	VTU-R	Definition
No peer xTU detected	4_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]
Any other or unknown initialization failure cause	5_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]
ITU-T G.998.4 retransmission mode was not selected while RTX_MODE = FORCED or with RTX_MODE = RTX_TESTMODE	6_{16}	M	M	See clause 7.5.1.6 of [ITU-T G.997.1]
Reserved for use by ITU-T	7_{16} to F_{16}			

11.5.1.4 Event type "Successful Init"

The event type "Successful Init" shall occur when the VTU transitions from the INIT/TRAIN to the SHOWTIME state (see VTU-O state diagram in Figure 12-2 and VTU-R state diagram in Figure 12-3). The related event data shall be the number of the bit to be set in xTU Transmission System Enabling (XTSE) as defined in clause 7.3.1.1.1 of [ITU-T G.997.1] for the representation of the transmission system type. For example, event data = 39_{16} (57 decimal) is ITU-T G.993.2 Region A (North America) (Annex A).

Table 11-47 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-47 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support").

Table 11-47 – Event data for event type "Successful Init"

Event data	ID	VTU-O	VTU-R	Definition
Reserved for use by ITU-T	00_{16}			
Transmission system type	01_{16} to 40_{16}	M	M	
Reserved for use by ITU-T	41_{16} to FF_{16}			
NOTE – This event may be used to correlate the timestamps in the VTU-O and VTU-R buffers.				

11.5.1.5 Event type "Downstream Init Net Data Rate"

The event type "Downstream Init Net Data Rate" shall occur when the VTU transitions from the INIT/TRAIN to the SHOWTIME state (see VTU-O state diagram in Figure 12-2 and VTU-R state diagram in Figure 12-3). The related event data shall identify the downstream net data rate applicable at the instant of entry into SHOWTIME.

The Downstream Init Net Data Rate (NDR) shall be recorded in the event data as an index, represented as an unsigned single octet integer, which shall be the integer value closest to:

$$63 \times \log_{10}\left(\frac{NDR}{0.2\ \text{Mbit/s}}\right) + 1, \text{ with } 1 \le \text{index} \le 253 \text{ for } 0.2\ \text{Mbit/s} \le \text{NDR} \le 2000\ \text{Mbit/s}.$$

If NDR≤ 0.2 Mbit/s, then the index shall be 0. If NDR ≥ 2000 Mbit/s, then index shall be 254. The index value 255 shall be reserved.

NOTE – This results in 3.7% granularity between recordable downstream net data rates.

11.5.1.6 Event type "Upstream Init Net Data Rate"

The event type "Upstream Init Net Data Rate" shall occur when the VTU transitions from the INIT/TRAIN to the SHOWTIME state (see VTU-O state diagram in Figure 12-2 and VTU-R state diagram in Figure 12-3). The related event data shall identify the upstream net data rate applicable at the instant of entry into SHOWTIME.

The Upstream Init Net Data Rate shall be recorded in the event data in the same way as the Downstream Init Net Data Rate (see clause 11.5.1.5).

11.5.1.7 Event type "Line Failure"

The event type "Line Failure" shall occur every time instant when one or more near-end line failures occurs or terminates (see Table 7-10 of [ITU-T G.997.1]). These line failures may or may not result in a retrain. The related event data shall consist of a bitmap identifying each of the line failures, which can simultaneously record multiple types of line failures. The event data indicates whether each line failure initially occurs or terminates in the reported second. If there are no line failures occurring or terminating in this second, then this event is not recorded.

Line failures are generated by the ITU-T G.997.1 functionality (see Figure 7-1 of [ITU-T G.997.1]). Therefore, the data gathering functionality gathers records from both the transceiver functionality and the ITU-T G.997.1 functionality.

Table 11-48 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-48 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support").

Table 11-48 – Event data for event type "Line Failure"

Event data	ID	VTU-O	VTU-R	Definition
LOS occurs	0 0 x x x x x 1	M	M	See clause 7.1.1.1.1 of [ITU-T G.997.1]
LOS terminates	0 0 x x x x 1 x	M	M	See clause 7.1.1.2.1 of [ITU-T G.997.1]
Loss of frame (LOF) occurs	0 0 x x x 1 x x	M	M	See clause 7.1.1.1.2 of [ITU-T G.997.1]
LOF terminates	0 0 x x 1 x x x	M	M	See clause 7.1.1.2.2 of [ITU-T G.997.1]
LPR occurs	0 0 x 1 x x x x	M	M	See clause 7.1.1.1.3 of [ITU-T G.997.1]
LPR terminates	0 0 1 x x x x x	M	M	See clause 7.1.1.2.3 of [ITU-T G.997.1]
NOTE – Event data denoted by "occurs" indicate the first occurrence of the event in a contiguous series.				

11.5.1.8 Event type "CRC-8 Anomalies"

The event type "CRC-8 Anomaly" records one or more contiguous seconds with CRC-8 anomalies at the near-end (see clause 9.5.2.3). The event data shall consist of a bitmap identifying the start and end of the CRC-8 anomalies at the near-end, which can simultaneously record multiple CRC-8 starting and ending events. For this event type; event data "1 or more CRC-8s anomalies per second started" shall occur at the first occurrence of a contiguous series of seconds with one or more CRC-8 anomalies; event data "1 or more CRC-8s anomalies per second ended" shall occur in the second immediately following the last occurrence of a contiguous series of seconds with one or more CRC-8 anomalies; event data "18 or more normalized CRC-8s anomalies per second started" shall occur at the first occurrence of a contiguous series of seconds with 18 or more normalized CRC-8 anomalies; and event data "18 or more normalized CRC-8s anomalies per second ended" shall occur in the second immediately following the last occurrence of a contiguous series of seconds with 18 or more normalized CRC-8 anomalies.

Table 11-49 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-49 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support").

Table 11-49 – Event data for event type "CRC-8 Anomalies"

Event data	ID	VTU-O	VTU-R	Definition
One or more CRC-8s anomalies per second starts	0 0 0 0 x x 0 1	M	M	See clause 9.5.2.3
One or more CRC-8s anomalies per second ended	0 0 0 0 x x 1 0	M	M	See clause 9.5.2.3
18 or more normalized CRC-8s anomalies per second starts	0 0 0 0 0 1 x x	M	M	See clause 9.5.2.3
18 or more normalized CRC-8s anomalies per second ended	0 0 0 0 1 0 x x	M	M	See clause 9.5.2.3
NOTE – One or more CRC-8 anomalies per second is one possible cause of an ES-L (see clause 7.2.1.1.2 of [ITU-T G.997.1]). 18 or more normalized CRC-8 anomalies per second is one possible cause of an SES-L (see clause 7.2.1.1.3 of [ITU-T G.997.1]).				

11.5.1.9 Event type "OLR"

The event type "OLR" shall occur when an (sent or received) OLR command is deferred, rejected or positively acknowledged (i.e., successful) (see clause 13.1) and when no response to a sent OLR command is received within the specified time period (i.e., failed) (see Table 11-1). The related event data shall identify the result of the OLR procedure. Table 11-50 lists the event data, and the identifier that shall be used in the record to identify the event data. If there are one or more events with a given event data in this second, then the indicator bit corresponding to that event data is set to 1; the indicator bit is set to 0 otherwise. Multiple instances of different event data for this event that occur during the same second shall be recorded by setting multiple indicator bits to 1 as indicated in Table 11-50.

Table 11-50 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support").

Table 11-50 – Event data for event type "OLR"

Event data	ID	VTU-O	VTU-R	Definition
Successful SRA downstream	x x x x x x x 1	M	M	See clause 13.1
Failed, deferred, or rejected SRA downstream	x x x x x x 1 x	M (see Note 1)	M	See clause 13.1
Successful SRA upstream	x x x x x 1 x x	M	M	See clause 13.1
Failed, deferred, or rejected SRA upstream	x x x x 1 x x x	M	M (see Note 1)	See clause 13.1
Successful SOS downstream	x x x 1 x x x x	M	M	See clause 13.1
Failed, deferred, or rejected SOS downstream	x x 1 x x x x x	M (see Note 1)	M	See clause 13.1
Successful SOS upstream	x 1 x x x x x x	M	M	See clause 13.1

Table 11-50 – Event data for event type "OLR"

Event data	ID	VTU-O	VTU-R	Definition
Failed, deferred, or rejected SOS upstream	1 x x x x x x x	M	M (see Note 1)	See clause 13.1
NOTE 1 – A failed event would not be recorded in this case. NOTE 2 – Event types that are optional in this Recommendation are mandatory for the data gathering function only if the particular event type is supported by the VTU.				

11.5.1.10 Event type "Bitswap"

The event type "Bitswap" shall occur when a bitswap command is deferred, rejected or positively acknowledged (i.e., successful), (see clause 13.1) and when no response to a sent bit swapping command is received within the specified time period (i.e., failed) (see Table 11-1). The related event data as per Table 11-51 shall identify the result of the bit swapping procedure.

If a particular bitswap event data occurs at least once, then the bit corresponding to the Bitswap event data in Table 11-51 shall be set to "1". If a particular bitswap event data did not occur, then the bit corresponding to the Bitswap event data in Table 11-51 shall be set to "0". The two highest MSB of the ID shall indicate the number of occurrences of this event during this second, i.e., one event, two events, and more than two events. Multiple instances of this event that occur in the same second shall be recorded simultaneously in a single record using the first two bits of the event data ID to record the total number of Bitswap events. Table 11-51 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-51 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support").

Table 11-51 – Event data for event type "Bitswap"

Event data	Two MSB of ID	VTU-O	VTU-R	Definition
Number of Bitswap events	0 1	M	M	1 = Number of Bitswap events
	1 0	M	M	2 = Number of Bitswap events
	1 1	M	M	3 = More than two Bitswap events
	Six LSB of ID	**VTU-O**	**VTU-R**	**Definition**
Event data				
Successful bitswap downstream	x x x x x 1	M	M	See clause 13.1
Failed, deferred or rejected bitswap downstream	x x x x 1 x	M (see Note)	M	See clause 13.1
Successful bitswap upstream	x x x 1 x x	M	M	See clause 13.1
Failed, deferred or rejected bitswap upstream	x x 1 x x x	M	M (see Note)	See clause 13.1
NOTE – A failed bitswap would not be recorded in this case.				

11.5.1.11 Event type "Downstream Net Data Rate after Successful SRA"

The event type "Downstream Net Data Rate after Successful SRA" shall occur when a downstream SRA request is positively acknowledged (i.e., successful). The related event data shall identify the downstream net data rate applicable at the completion of the SRA procedure.

The Downstream Net Data Rate after Successful SRA shall be recorded in the event data in the same way as the Downstream Init Net Data Rate (see clause 11.5.1.5).

11.5.1.12 Event type "Upstream Net Data Rate after Successful SRA"

The event type "Upstream Net Data Rate after Successful SRA" shall occur when an upstream SRA request is positively acknowledged (i.e., successful). The related event data shall identify the upstream net data rate applicable at the completion of the SRA procedure.

The Upstream Net Data Rate after Successful SRA shall be recorded in the event data in the same way as the Downstream Init Net Data Rate (see clause 11.5.1.5).

11.5.1.13 Event type "Downstream Net Data Rate after Successful SOS"

The event type "Downstream Net Data Rate after Successful SOS" shall occur when a downstream SOS request is positively acknowledged (i.e., successful). The related event data shall identify the downstream net data rate applicable at the completion of the SOS procedure.

The Downstream Net Data Rate after Successful SOS shall be recorded in the event data in the same way as the Downstream Init Net Data Rate (see clause 11.5.1.5).

11.5.1.14 Event type "Upstream Net Data Rate after Successful SOS"

The event type "Upstream Net Data Rate after Successful SOS" shall occur when an upstream SOS request is positively acknowledged. The related event data shall identify the upstream net data rate applicable at the completion of the SOS procedure.

The Upstream Net Data Rate after Successful SOS shall be recorded in the event data in the same way as the Downstream Init Net Data Rate (see clause 11.5.1.5).

11.5.1.15 Event type "Defect"

If a *los*, *lom*, or *sef* defect (see clauses 11.3.1.3 and 11.3.1.4) occurs during any part of a second, then that second is considered to be a *los* defected second, a *lom* defected second, or a *sef* defected second, respectively. The event type "Defect" shall occur each second when any type of defected second starts or ended. A type of defected second shall be recorded as "starts" when that type of defected second occurs in the current second but did not occur in the previous second. A type of defected second shall be recorded as having "ended" when that type of defected second occurred in the previous second but does not occur in the current second. A particular type of defected second (*los* defected second, *lom* defected second, *or sef* defected second) cannot both start and stop in the same second. In order to gather multiple types of defected second events in one record, the event data shall consist of a bitmap identifying the start and end each of each type of defected second.

Table 11-52 lists the event data and the identifier that shall be used in the event record to identify the event data. Table 11-52 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support" and "NA" indicates "not applicable").

Table 11-52 – Event data for event type "Defect"

Event data	ID	VTU-O	VTU-R	Definition
los second starts	x x x x x x 0 1	M	M	See clause 11.3.1.3
los second ended	x x x x x x 1 0	M	M	See clause 11.3.1.3
lom second starts	x x x x 0 1 x x	M	M	See clause 11.3.1.3
lom second ended	x x x x 1 0 x x	M	M	See clause 11.3.1.3
sef second starts	x x 0 1 x x x x	M	M	See clause 11.3.1.3
sef second ended	x x 1 0 x x x x	M	M	See clause 11.3.1.3

11.5.1.16 Event type "Retransmission Event"

The event type "Retransmission Event" shall occur when the VTU operates in ITU-T G.998.4 (retransmission) mode and one of the retransmission events (see Table 11-53) occurs. The related event data shall identify the retransmission event that occurred.

Table 11-53 lists the event data, and the identifier that shall be used in the record to identify the event data. Table 11-53 also indicates whether the event data shall be supported by the VTU-O and/or VTU-R (M indicates "mandatory support").

Table 11-53 – Event data for event type "Retransmission Event"

Event data	ID	VTU-O	VTU-R	Definition
Reserved for future use	00_{16}			
leftr defect	01_{16}	M	M	See clause 11.3.3 of [ITU-T G.998.4]
Reserved for use by ITU-T	02_{16} to FF_{16}			
NOTE – The *leftr* defect is recorded only for upstream at the VTU-O, and it is reported only for downstream at the VTU-R. This is because the *leftr* defect is near-end only.				

11.5.2 Buffer management rules

In order to obtain a fair mixing of event types in the buffer, while guaranteeing that the oldest entries per event type are removed first, the following buffer management rules shall apply:

- The sum of the configured logging depths, $logging\_depth\_event_i$, over all event types shall be smaller than or equal to the number of storable records in the buffer, *logging_depth*;
- When the buffer is not yet full, a new entry can be added without constraints;
- When the buffer is full, before adding a new entry, one entry shall be removed. The entry that shall be removed is the oldest entry out of all event types which are already above their configured logging depths or which would get above their configured logging depths if the new entry was to be added without a removal;

  NOTE – The third rule is described as a "remove before add". It can also be rephrased as "remove after add" as follows: "when the buffer is full, the new entry is added to a spare location. The entry that shall be removed is the oldest entry out of all event types which are above their configured logging depths."

These buffer management rules are illustrated in Figure 11-6.

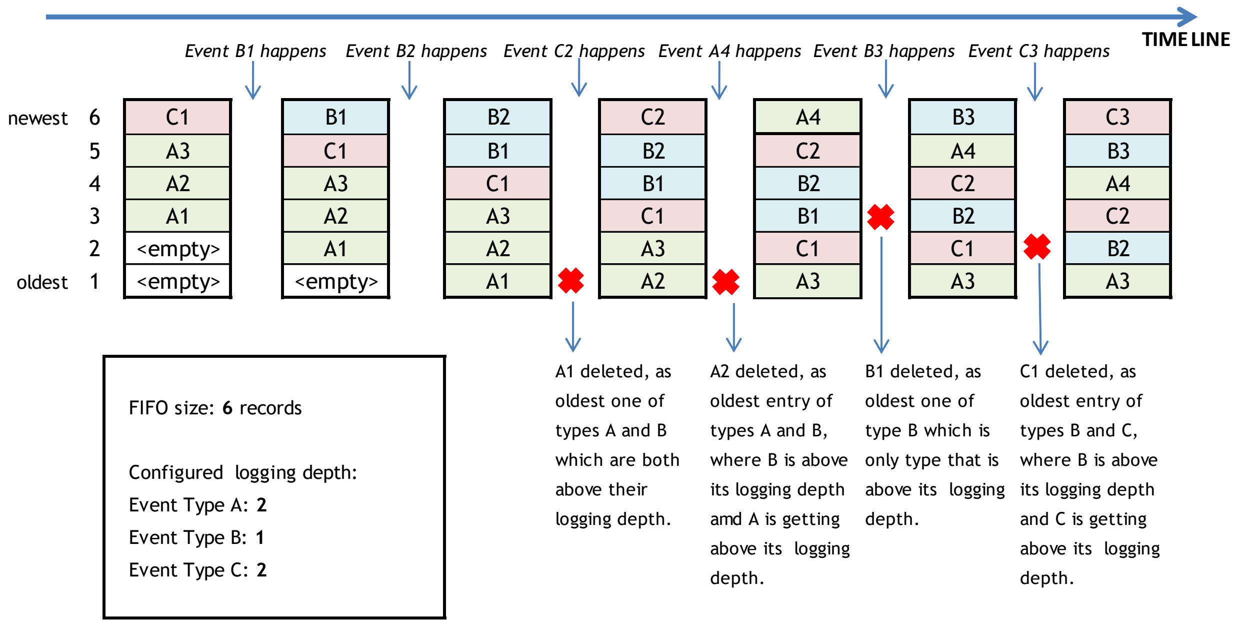

Figure 11-6 – Illustration of the buffer management rules

12 Link activation methods and procedures

12.1 Overview

12.1.1 Link states and link state diagram

The VDSL2 link states and activation/deactivation procedures diagram is illustrated in Figure 12-1.

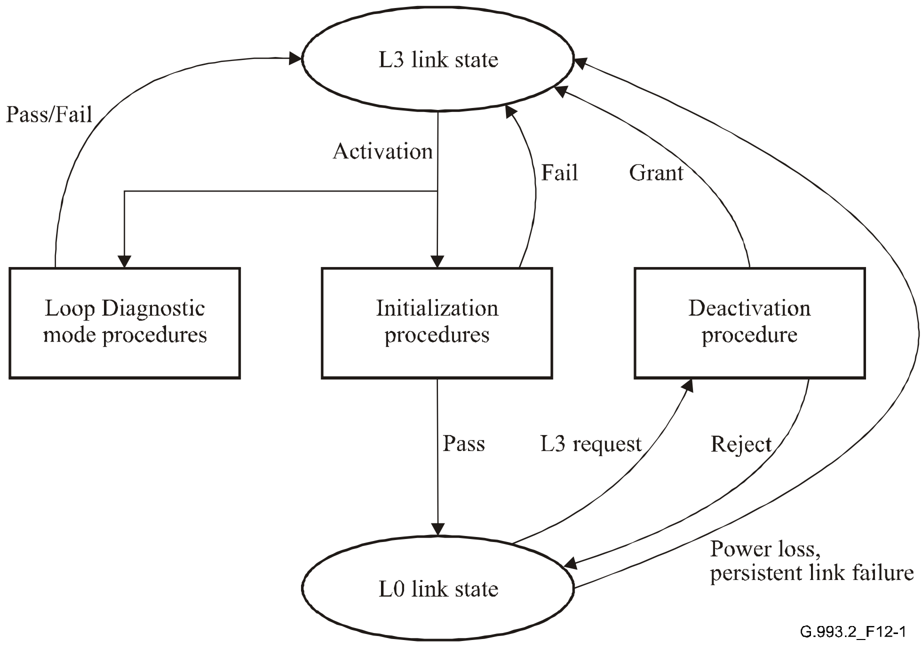

Figure 12-1 – VDSL2 link states and link state diagram

Figure 12-1 has two link states (L0 and L3), and also contains the procedures that allow the link to change from one link state to another. The link states are shown in rounded boxes, whilst the procedures are shown as rectangular boxes.

L3 is the link state where the VTU is provisioned through a management interface for the service desired by the operator. In this link state, both the VTU-O and VTU-R do not transmit any signal.

L0 is the link state achieved after the initialization procedure has completed successfully by both VTUs. In this link state, the link shall transport user information with standard performance characteristics according to the CO-MIB configuration.

12.1.2 Transceiver states and transceiver state diagram

State diagrams are given in Figure 12-2 for the VTU-O, and in Figure 12-3 for the VTU-R. States are indicated by ovals, with the name of the state given within the oval. The states are defined in Table 12-1 for the VTU-O and in Table 12-2 for the VTU-R. Transitions between states are indicated by arrows, with the event causing the transition listed next to the arrow. All states are mandatory.

A variety of "host controller" commands (events preceded by "c:_" and "r:_") are shown as non-mandatory in either state diagram to provide example events and transitions between states. The way in which these events are implemented is left to the vendor since many options are possible.

In the state diagram for the VTU-O, a C-IDLE state would be desired to guarantee a quiet mode, which may be useful to allow certain tests (e.g., MLT), or to discontinue service.

In the state diagram for the VTU-R, a self-test function is desirable, but it may be a vendor/customer option to define when self-test occurs (e.g., always at power-up or only under VTU-O control), and which transition to take after successfully completing self-test (e.g., enter R-IDLE, or enter R-SILENT).

IDLE is the state where the VTU is provisioned through a management interface for the service desired by the operator. In this state, the VTU does not transmit any signal. A VTU that receives a higher layer signal to activate (c:_L0_request for VTU-O or r:_L0_request for VTU-R) shall use the initialization procedure defined in clause 12.3 to transition the link from the L3 to the L0 state. A VTU that detects the signals of the initialization procedure at the U reference point, if enabled, shall respond by using the initialization procedure. If disabled, the VTU shall remain in the IDLE state.

The link transitions to the L0 state once the initialization procedure has completed successfully and both VTUs are in the SHOWTIME state. A VTU-O shall return to the O-SILENT state upon a guided power management (c:_L3_request, see clause 11.2.3.9), or upon a re-initialization triggered by Re-Initialization Policy (see clause 12.1.4). A VTU-R shall return to the R-SILENT state upon a guided power management (r:_L3_request, see clause 11.2.3.9), or upon a re-initialization triggered by Re-Initialization Policy (see clause 12.1.4). With the former, a VTU-R shall set AUTO_init=OFF to disable autonomous proceeding to the R-INIT/HS state. With the latter, a VTU-R shall set AUTO_init=ON to enable autonomous proceeding to the R-INIT/HS state.

The receiving VTU shall transition state upon Persistent LOS and/or LOF failure (see clause 12.1.4). This implies that if no high_BER-hs or high_BER-fs events cause the receiving VTU to transition state earlier, then the persistency allows the transmitting VTU to detect the LOS or LOF failure condition through the indicator bits, before the receiving VTU transitions state (i.e., removes the Showtime signal from the line).

NOTE – High_BER-fs event relates to fast start-up, which is for further study (see clause 12.5).

The receiving VTU shall also transition state upon a high_BER event (see clause 12.1.4). This event relates to near-end and/or far-end performance primitives and performance counters for which thresholds may be configured through the CO-MIB as to declare a high_BER event upon threshold crossing.

If the VTU-O transitions from O-SHOWTIME to O-SILENT, then the VTU-R shall detect a Persistent LOS Failure, shall transition to R-SILENT followed by R-INIT/HS and shall transmit R-TONES-REQ within a maximum of 6 s after the VTU-O transitioning to O-SILENT.

If the VTU-R transitions from R-SHOWTIME to R-SILENT, then the VTU-O shall detect a Persistent LOS Failure, shall transition to O-SILENT, either followed by waiting to receive R-TONES-REQ (VTU-R initiated HS) or followed by O-INIT/HS (VTU-O initiated HS).

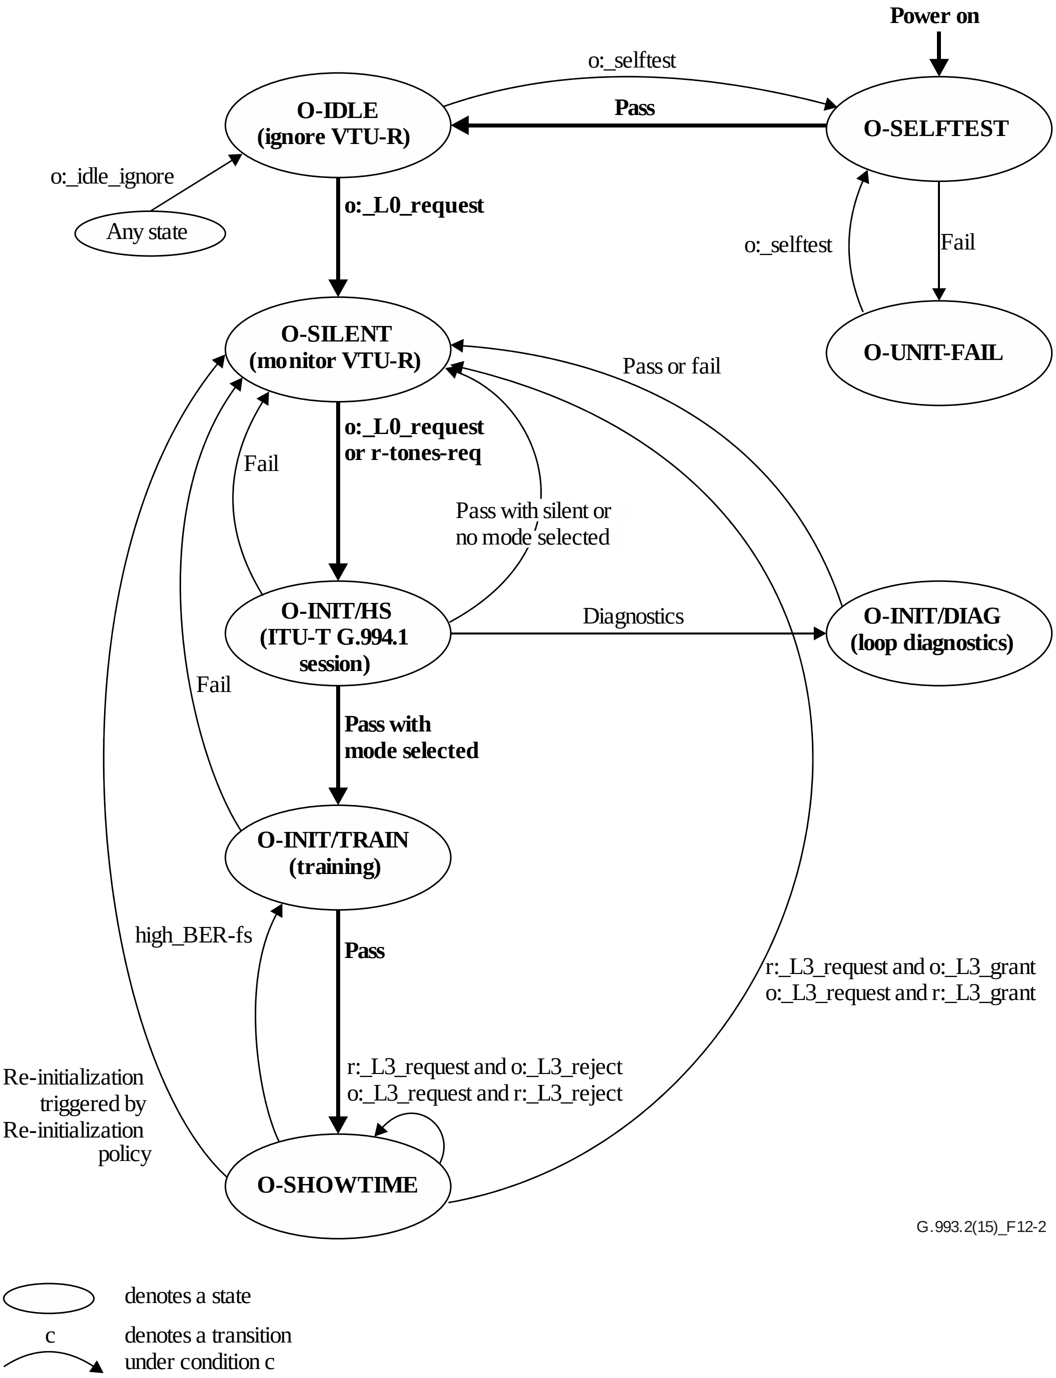

Figure 12-2 – State diagram for the VTU-O

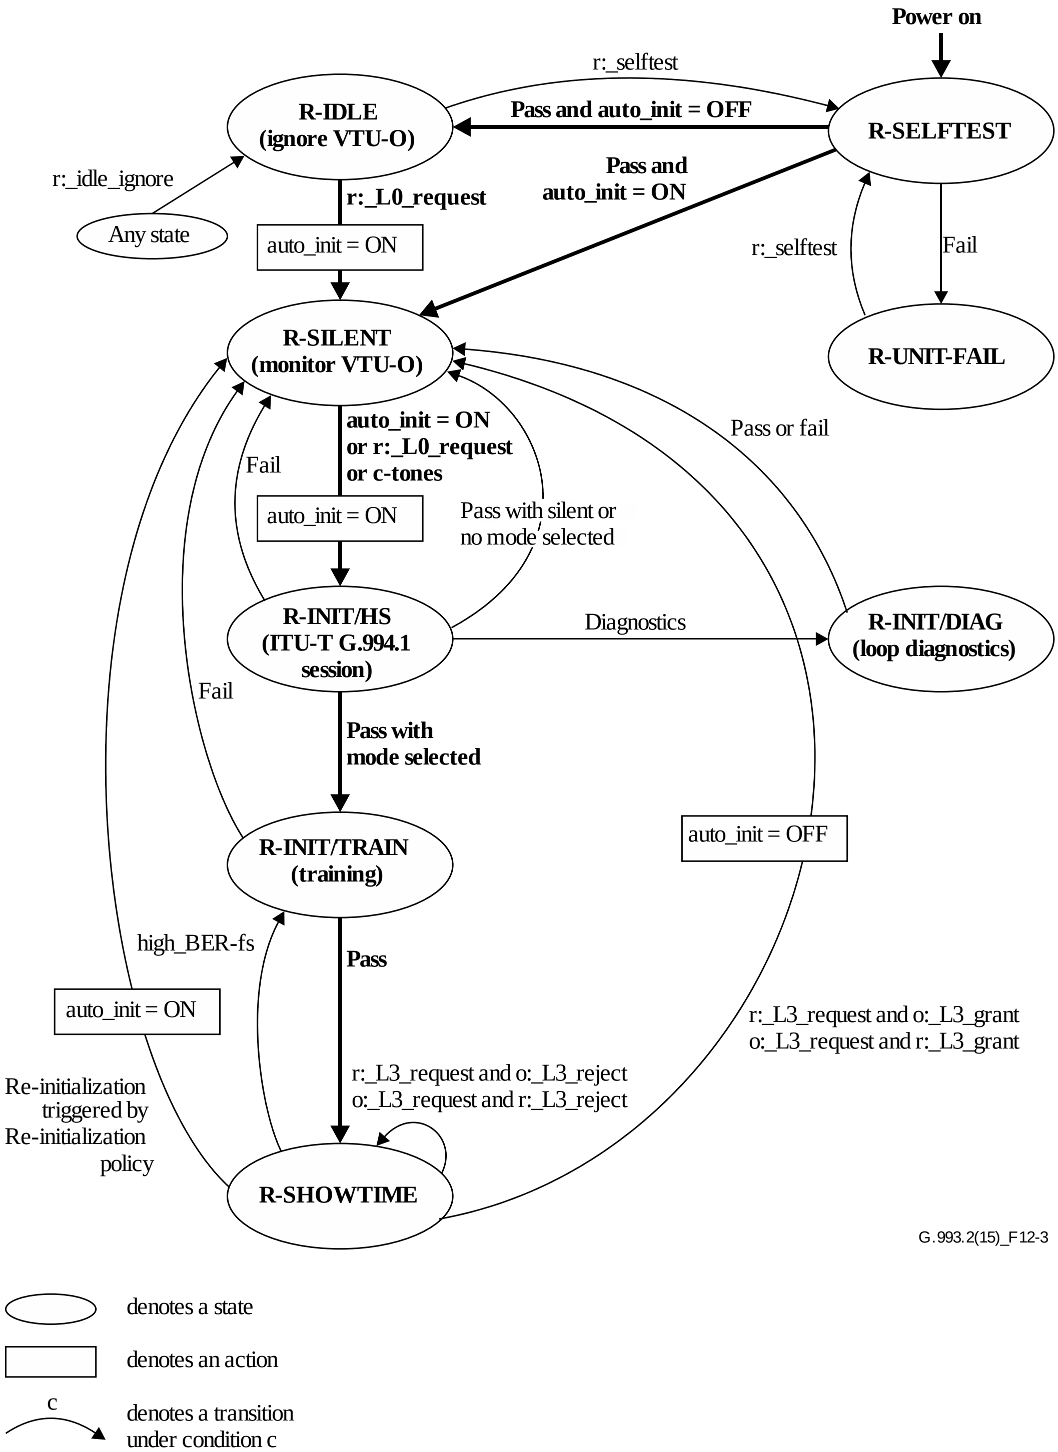

Figure 12-3 – State diagram for the VTU-R

Table 12-1 – VTU-O state definitions

State name	Description
O-SELFTEST (mandatory)	• Temporary state entered after power-up in which the VTU performs a self-test; • Transmitter off (QUIET at U-O interface); • Receiver off (no response to R-TONES-REQ signal); • No response to host control channel; • If self-test pass then transition to O-IDLE; • If self-test fail then transition to O-UNIT-FAIL.
O-UNIT-FAIL (mandatory)	• Steady state entered after an unsuccessful VTU self-test; • Transmitter off (QUIET at U-O interface); • Receiver off (no response to R-TONES-REQ signal); • Monitor host control channel if possible (allows the host controller to retrieve self-test results).
O-IDLE (mandatory)	• Steady state entered after successful self-test; • Transmitter off (QUIET at U-O interface); • Receiver off (no response to R-TONES-REQ signal); • Monitor host control channel.
O-SILENT (mandatory)	• Steady state defined in ITU-T G.994.1, entered upon host controller command; • Transmitter off (QUIET at U-O interface); • Receiver on (monitor for R-TONES-REQ signal, if detected, transition to O-INIT/HS state); • Monitor host control channel.
O-INIT/HS (mandatory)	• Temporary state entered to perform ITU-T G.994.1 phase of initialization; • Transmitter on (start with transmitting C-TONES signal); • Receiver on (start with monitoring for R-SILENT0 signal); • Monitor host control channel; • If silent period or no mode selected then transition to O-SILENT1; • If loop diagnostics mode then transition to O-INIT/DIAG; • If operating mode selected then transition to O-INIT/TRAIN.
O-INIT/TRAIN (mandatory)	• Temporary state entered to perform other phases of initialization; • Transmitter on (start with O-P-QUIET1); • Receiver on (start with monitoring for R-P-QUIET1); • If init pass then transition to O-SHOWTIME; • If init fail then transition to O-SILENT; • Monitor host control channel.
O-INIT/DIAG (mandatory)	• Temporary state entered to perform other phases of initialization in loop diagnostics mode; • Transmitter on (start with O-P-QUIET1); • Receiver on (start with monitoring for R-P-QUIET1); • Transition to O-SILENT; • Monitor host control channel.
O-SHOWTIME (mandatory)	• Steady state entered to perform bit pump functions (frame bearers active); • On-line reconfigurations occur within this state; • Upon conditions satisfying the Re-Initialization Policy ($RIpolicy_n$) then transition to O-SILENT; • If link transition to L3 state is granted, then transition to O-SILENT; • Monitor host control channel.

Table 12-2 – VTU-R state definitions

State name	Description
R-SELFTEST (mandatory)	• Temporary state entered after power-up in which the VTU performs a self-test; • Transmitter off (QUIET at U-R interface); • Receiver off (no response to C-TONES signal); • No response to host control channel; • If self-test pass then transition to R-IDLE if VTU is under host control or transition to R-SILENT if VTU is in automatic training mode; • If self-test fail then transition to R-UNIT-FAIL.
R-UNIT-FAIL (mandatory)	• Steady state entered after an unsuccessful VTU self-test; • Transmitter off (QUIET at U-R interface); • Receiver off (no response to C-TONES signal); • Monitor host control channel if possible (allows the host controller to retrieve self-test results).
R-IDLE (mandatory)	• Steady state entered after successful self-test if VTU is under host control; • Transmitter off (QUIET at U-R interface); • Receiver off (no response to C-TONES signal); • Monitor host control channel.
R-SILENT (mandatory)	• Temporary state defined in[ITU-T G.994.1] entered after self-test pass if VTU is in automatic training mode or with host controller command; • Transmitter off (transmit R-SILENT0 signal); • Receiver on (monitor for C-TONES signal, if detected, transition to R-INIT/HS state); • Automatic training: immediate transition to R-INIT/HS (unless delayed for silent period or in orderly shutdown condition); • Monitor host control channel.
R-INIT/HS (mandatory)	• Temporary state entered to perform ITU-T G.994.1 phase of initialization; • Transmitter on (start with transmitting R-TONES-REQ signal); • Receiver on (start with monitoring for C-TONES signal); • Monitor host control channel; • If silent period or no mode selected then transition to R-SILENT; • If loop diagnostics mode then transition to R-INIT/DIAG; • If operating mode selected then transition to R-INIT/TRAIN.
R-INIT/TRAIN (mandatory)	• Temporary state entered to perform other phases of initialization; • Transmitter on (start with R-P-QUIET1 signal); • Receiver on (start with monitoring for C-P-QUIET1 signal); • If init pass then transition to R-SHOWTIME; • If init fail then transition to R-SILENT; • Monitor host control channel.
R-INIT/DIAG (mandatory)	• Temporary state entered to perform other phases of initialization in loop diagnostics mode; • Transmitter on (start with R-P-QUIET1); • Receiver on (start with monitoring for C-P-QUIET1); • Transition to R-SILENT; • Monitor host control channel.

Table 12-2 – VTU-R state definitions

State name	Description
R-SHOWTIME (mandatory)	• Steady state entered to perform bit pump functions (frame bearers active); • On-line reconfigurations occur within this state; • Upon conditions satisfying the Re-Initialization Policy ($RIpolicy_n$) then transition to R-SILENT; • If link transition to L3 state is granted, then transition to R-SILENT; • Monitor host control channel.

12.1.3 Initialization procedures

During the ITU-T G.994.1 handshake phase of the initialization procedure, the VTUs exchange capability lists and agree on a common mode for training and operation using the ITU-T G.994.1 protocol. A successful completion of the ITU-T G.994.1 handshake phase will lead to either the channel discovery phase of initialization or the loop diagnostic mode (depending on which one is selected). Failure of the ITU-T G.994.1 handshake phase leads the VTU back to the SILENT state and leads the link back to the L3 state. The handshake procedure is described in clause 12.3.2 and [ITU-T G.994.1].

During the channel discovery, training, and channel analysis and exchange phases of initialization, the VTUs train their respective transceivers after identifying the common mode of operation. During these phases, the transceivers identify channel conditions, exchange parameters for Showtime operation, etc. After successful completion of the initialization procedure, the transceivers transition to the SHOWTIME state (Showtime). Upon unsuccessful completion of the initialization procedure, the VTUs return to the SILENT state and the link returns to the L3 state. The initialization phases are described in clauses 12.3.3 through 12.3.5.

12.1.4 Deactivation, power loss, persistent link failure and high_BER events

The deactivation procedure allows an orderly shutdown of the link. The VTUs shall follow the procedures described in clause 11.2.3.9 to transition the link from the L0 state to the L3 state.

The link is in the L3 state, after both VTU-O and VTU-R have transitioned from the SHOWTIME state to the SILENT state.

Two policies are defined for the VTU to trigger a transition from the SHOWTIME state to the SILENT state. The selection of the policy is controlled via the parameter "Re-Initialization Policy" ($RIpolicy_n$).

In the first policy ($RIpolicy_n$=0) (mandatory), a VTU shall transition from the SHOWTIME state to the SILENT state in the case of:

1. loss of receive power (power loss); or
2. persistent link failure; or
3. upon a high_BER-hs event as defined below for $RIpolicy_n$=0.

The VTU shall declare a power loss when a persistent LOS failure is declared. Persistent LOS failure is declared after 2.5 ± 0.5 s of near-end LOS failure with the *los* (see clause 11.3.1.3) still present. An LOS failure is declared after 2.5 ± 0.5 s of contiguous *los*, or, if *los* is present when the criteria for LOF failure declaration have been met (see LOF Failure definition below). An LOS failure is cleared after 10 ± 0.5 s of no *los*.

The VTU shall declare a persistent link failure when a persistent LOF failure is declared. A persistent LOF failure is declared after 2.5 ± 0.5 s of near-end LOF failure with the *sef* (see clause 11.3.1.3) still present. An LOF failure is declared after 2.5 ± 0.5 s of contiguous near-end *sef*, except when a *los* or