




White Paper
IPTV QoE:
Understanding and
interpreting MDI values



This paper gives an overview of the MDI and discusses some potential applications related to measuring video quality.



Agilent Technologies

Understanding the Media Delivery Index

As service providers actively deploy triple play networks that deliver voice, video and data over a converged infrastructure, network equipment manufacturers are fervently designing and testing devices that enable quality of service (QoS) on these networks. QoS mechanisms allow devices to apply policies to the different types of traffic present on a network in order to ensure that each one is treated in the most appropriate way. Voice traffic, for example, typically receives high priority because it is very sensitive to delay. Data traffic, such as web or email, is not as affected by timing and therefore does not need the same preferential treatment.

Naturally, consumers are not concerned with traffic priority and dropped packets: they want their phone calls to be clear and their IPTV programs to be smooth and free from visual impairments. From this perspective, it is the users' quality of experience (QoE) that really matters. Fickle customers in this highly competitive landscape will, without compunction, dump providers who do not meet their expectations of quality. Hence, in order to be successful, service providers and network equipment manufacturers must rigorously test their triple play devices and ensure that they provide a proper QoE.

The video component of the triple play offering presents unique demands on the network because of its high bandwidth requirements and low tolerance to jitter and packet loss. The media delivery index (MDI) measurement gives an indication of expected video quality – ultimately, users' QoE – based on network level measurements. It is independent of the video encoding scheme and is a lightweight, scalable alternative to measurements such as MPQM and V-Factor that decode and examine the video itself. This paper gives an overview of the MDI and recommended acceptable measurement values, and discusses some potential applications related to measuring video quality.

Components of the Media Delivery Index (MDI)

A typical network infrastructure is illustrated in Fig 1. The media delivery index measurements are cumulative throughout the network and can be measured from any point between the video sources and set top boxes (STBs). The MDI is typically displayed as two numbers separated by a colon: the delay factor (DF) and the media loss rate (MLR).

DF:MLR

Delay Factor (DF)

In order to understand the delay factor component of the MDI, it is useful to revisit the relationship between jitter and buffering. Jitter is a change in end-to-end latency with respect to time. Packets arriving at a destination at a constant rate exhibit zero jitter. Packets with an irregular arrival rate exhibit non-zero jitter. Fig. 2 illustrates this difference.




Fig. 2: Packets received versus time with and without jitter.




Fig. 1: A typical network infrastructure

After traversing the network separating the data source and the destination, and being queued, routed and switched by the various intervening network elements, packets are liable to arrive at the destination with some rate variation over time. This can, for example, be caused by transient network congestion due to a large amount *peer-to-peer* (P2P) traffic, dynamic subscriber actions such as placing a VoIP call that preempts the video traffic, or to packets taking different paths through the IP network. In any event, if the instantaneous data arrival rate does not match the rate at which the destination is consuming data, the packets must be buffered upon arrival.

Consider a typical 3.75 Mb/s MPEG video transport stream. The decoder at the destination will consume (or drain) a constant 3.75 Mb/s of data, but the data may arrive at rates above or below the drain rate. Buffers in the decoder are used to collect a certain number of packets, arriving at different rates, and feed them to the decoding engine at a constant rate.

The more severe the jitter, the larger the buffers need to be in order to eliminate it. The price to pay for having larger buffers is that they introduce delay. Furthermore, buffers are of a finite size, and excessive jitter will cause them to either overflow or underflow. An overflow is when packets are arriving at such a high rate that they fill the buffer and cause packets to be dropped at the receiver.

An underflow is when packets are arriving so slowly that the buffers do not have enough data to feed the decoder at its drain rate. Both of these situations are undesirable, and they degrade the user's QoE. Users may see the video pause or play in a choppy fashion and the images can contain visual distortions as a result of the lost media packets.

The DF component of the MDI is a time value indicating how many milliseconds' worth of data the buffers must be able to contain in order to eliminate jitter. It is computed as packets arrive and is displayed to the user at regular intervals (typically one second). It is calculated as follows:

1. At every packet arrival, calculate the difference between the bytes received and the bytes drained. This is the MDI virtual buffer depth Δ

$$\Delta = |\text{bytes_received} - \text{bytes_drained}|$$

2. Over a time interval, take the difference between the minimum and maximum values of Δ and divide by the media rate:

$$DF = \frac{(\max(\Delta) - \min(\Delta))}{\text{media_rate}}$$

As an example, consider once again a 3.75Mb/s MPEG video stream. If, over the course of the one-second interval, the maximum amount of data in the virtual buffer is 3.755 Mb and the minimum amount is 3.740 Mb, the delay factor would be computed as:

$$DF = \frac{3.755\text{Mb} - 3.740\text{Mb}}{3.75\text{Mb/s}} = \frac{15\text{kb}}{3.75\text{Mb/s}} = 4\text{ms}$$

Hence, in order to avoid packet loss in the presence of the observed jitter, the receiver's buffer would have to be of 15kb, which would inject 4 milliseconds of delay.

The DF can be employed at the video destination to assess video quality from the user's perspective, and infer QoE. It can also be used to determine the impact of each network element along the video delivery path. By comparing the DF at the ingress of a device to the DF at the egress, it is possible to determine the device's footprint. Devices that do not inject jitter will have a smaller footprint, and are better suited to delivering video.

Acceptable Delay Factor (DF)

Maximum Acceptable DF
9-50 ms

Fig. 3: Recommended maximum acceptable DF

The delay factor that is acceptable for any particular network varies greatly because of the wide range of buffer sizes available. Most STBs use a single RAM module. Only part of this RAM is actually used as a buffer for de-jittering incoming IP streams, which is why most STB specifications do not list the buffer size. The actual buffer size of each STB can thus only be determined through testing with specialized hardware. Triple play QoE standards that are being developed in the DSL Forum as part of WT-126, recommend that the jitter introduced within the network remain below 50 ms, however, this is much less than what mid to high-end STBs can actually handle. Tests performed at Agilent determined the maximum acceptable DF for a particular low-end STB to be 9 ms. These two values vary slightly depending on the stream rate of the codec, but the amount is negligible (less than 10%). The discrepancy between the two values (50 ms and 9 ms) is attributed to a wide variation in the quality of available STBs, thus the recommendations are for different STB buffer sizes. The exact maximum acceptable DF must be tuned to the buffer size of the STBs; to do this find the highest amount of jitter your STB can handle before any visual distortions appear (maintain zero packet loss).

Media Loss Rate (MLR)

The *media loss rate* is simply defined as the number of lost or out-of-order media packets per second. Out-of-order packets are important because many devices make no attempt to reorder packets before presenting them to the decoder. Any packet loss – represented as a non-zero MLR – will adversely affect video quality and can introduce visual distortions or irregular, uneven video playback. MLR is a convenient format for specifying *service level agreements* (SLAs) in terms of packet loss rates. So, taken in context with the previous DF component, a device with an MDI of 4:0.001 would indicate that the device has a delay factor of 4 milliseconds and a media loss rate of 0.001 media packet per second.

Acceptable Media Loss Rate (MLR)

Maximum Acceptable Channel Zapping MLR
0

Fig. 4: Recommended maximum acceptable channel zapping MLR for all services and codecs

Service (All Codecs)	MaxAcceptable Average MLR
SDTV	0.004
VOD	0.004
HDTV	0.0005

Fig. 5: Recommended maximum acceptable average MLRs

Because the *media loss rate* is a rate, some important information is lost, such as whether the IP packets lost are consecutive or inconsecutive. A study has been performed as part of the research into the DSL Forum's WT-126 where it has been shown that almost all single IP packets lost produced a visible error, and a typical user prefers less frequent but significant errors over more frequent but less significant errors. QoE standards for IPTV are still under debate, but WT-126 currently recommends a maximum loss of up to five consecutive IP packets per thirty minutes for SDTV and VOD, and four hours for HDTV. If translated into MLR terms, this assumes the loss is a single IP packet in the specified timeframe. To understand why, assume the MLR is based on five IP packets lost in the specified timeframe; this would mean a maximum acceptable MLR of 0.019 (assuming seven media packets per IP packet):

$$\frac{5 \text{ IP Packets}}{30 \text{ minute}} \times \frac{1 \text{ minute}}{60 \text{ second}} \times \frac{7 \text{ Media Packets}}{1 \text{ IP Packet}} = \frac{0.019 \text{ Media Packets}}{\text{second}}$$

This rate implies that it is acceptable to have five inconsecutive loss events (each event consisting of a single IP packet) per thirty minutes, which is not the case. This is why the loss scenario considered is a single IP packet or event.

The maximum acceptable MLR also depends on the implementation. For channel zapping, a channel is generally viewed for a brief period, and thus one would be interested if any packet loss at all occurred. For this case the maximum acceptable MLR is 0, as stated in Fig. 4, because any greater a value would mean a loss of 1 or more packets in a small viewing timeframe; much more than what is acceptable following current QoE standards.

The second implementation is network monitoring, in which the sample period is generally large. This means the number of packets lost is best expressed as an average per second rather than a total, as illustrated in Fig. 5

Applications of the MDI

The MDI is useful for locating and characterizing network issues that can adversely affect media quality and users' QoE. If the MDI is tracked at intermediate points in the delivery network, the difference in the DF and MLR components between successive network elements can help quickly isolate the source of potential or actual impairments. If a large MLR is recorded at one router whilst the MLR was zero at the previous hop in a stream's path, it is a strong indication that something unfavorable occurred in that network segment, such as a buffer overflow or packet corruption. Similarly, if the delay factor DF jumps significantly between two successive hops, this may suggest long queuing delays due to congestion. It is also a warning of impending packet loss. As discussed, with larger buffers, it is possible to compensate for high jitter, at the cost of added delay. An alternative interpretation of the MDI is that, through the DF, it characterizes how much margin (buffer space) is required before media quality is adversely impacted.

MDI measurements can also be used in lab environments to determine the effect that particular devices may have on media quality. Injecting a large number of streams with known MDI characteristics and observing the output MDI of the streams is one useful way to characterize a device's suitability for delivering video. It is also possible to observe the effect of various network events, such as congestion or routing protocol updates, on users' QoE by making MDI measurements in the presence of different control and data traffic stimuli.

Conclusion

In summary, the media delivery index is a lightweight, scalable metric for assessing the effect a delivery network has on video and, ultimately, on end user QoE. Its two components, the delay factor and media loss rate, use packet loss and jitter as predictors of IPTV quality. Network infrastructure testing using measurements such as MDI is vital to the success of IPTV deployments: customer loyalty can only be secured with a good quality service, which, in turn, can only be ensured with rigorous testing.

IPTV Test Solutions from Agilent Technologies

Agilent Technologies is the world's premier measurement company. Agilent delivers test solutions that provide rapid insight and accelerate time to market for its customers as they develop and deploy devices for Triple Play networks. Agilent's N2X platform of high-performance test systems provides a complete IPTV test solution, including highly scalable video generation and analysis capabilities. More information can be found on the web at www.agilent.com/comms/N2X, or by contacting your local Agilent salesperson or distributor.

This page intentionally left blank.

This page intentionally left blank.

Sales, Service and Support

United States:

Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:

Agilent Technologies Canada Inc.
2660 Matheson Blvd. E
Mississauga, Ontario
L4V 5M2
1-877-894-4414

Europe:

Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom

07004 666666

Japan:

Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:

Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 269-7500
Fax: (305) 267-4286

Asia Pacific:

Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
Fax: (852) 2506-9233

Australia/New Zealand:

Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

This information is subject to change without notice.

Printed on recycled paper

© Agilent Technologies, Inc. 2006

Printed in USA August 30, 2006

5989-5088EN



Agilent Technologies