




Delay and Latency




White Paper

Some customers make an issue of delay and latency while others seem not to care. Some vendors also claim “low latency” while others are silent on the issue. Why is this? What is the deal with delay?

First of all, let’s agree that “delay” and “latency” are the same thing. Officially, we think of “delay” as the perception of a technical artifact called “latency”, but let’s just call the whole thing “delay”.

We define delay to be the amount of time from it takes for light hitting a camera that is connected to one VBrick encoder to be transmitted, decoded, and displayed on a TV screen that is connected to another VBrick. In other words, delay is the time between “light in to light out”. We assume audio/video sync, so the same can be said about audio too.

The best case would be if delay were zero, but it never is. At the very least, you have what’s called “propagation delay”, which is the time it takes for signals to propagate over wires. Okay, this is about the speed of light, but remember that we are sometimes talking about very long distances. For example, a satellite in geosynchronous orbit 22,000 miles (36,000 Km) up will introduce between 239.6 ms and 279.0 ms in each direction. So that’s about half a second total delay just in the satellite system even though radio travels at the speed of light!



As I sit here in Wallingford Connecticut USA, I can ping remote computers on the Internet. I know that it takes about 125 ms to reach a web site hosted 3,000 miles away in California. Most of that delay has to do with the many Internet router points in between.

The point is that there is some minimum delay that will always be present. For high speed local networks, the propagation delay is very low and can be ignored.

The question is: what total delay is ‘good enough’? For one-way streaming, the answer could actually be “don’t care” and several minutes would be acceptable. Windows Media, for example, has a delay of about 5 to 15 seconds which is perfectly acceptable for most web streaming applications.

But what is a reasonable delay for a video conference? The answer is 300 ms. This is not a VBrick answer, but one that is generally accepted in the industry. Beyond 300 ms, you cannot have an interactive session, and find yourself needing to say “over” at the end of each sentence!

Video Delay

You know that NTSC video runs at 30 fps, and PAL at 25 fps. In order to send even one frame of video, it has to be digitized. Well, you can’t digitize it until you’ve received it, right? So, you can assume there has to be at least a 1 frame delay (that’s 33 ms). MJPEG, which provides no inter-frame compression, compresses each frame and can introduce as little as one frame of

delay...in practice it is usually more than this, but still quite low. Of course, MJPEG does not provide the compression that you get with MPEG and requires vast amounts of bandwidth to produce acceptable video.

MPEG

MPEG too has at least one frame of delay. But MPEG can generate "B" frames ("Bi-directional Frames") which can only be calculated using several frames of video. B frames, therefore, create a lot of delay. Interestingly, it's not just an encoder question: a MPEG encoder can create I frames only and behave somewhat like MJPEG, and it can create both I frames then P frames. It can also calculate B frames to provide higher video quality at a given data rate, and the I-B-P Group of Pictures is an important part of MPEG. But an MPEG encoder generates first I frames, then P frames, and only then can it calculate B frames. And because the B frames "belong" between the I and P frames, a MPEG decoder must buffer the data and re-order it. B frames is one cause for the delay sometimes found in MPEG systems.

As a doctor would say, if it hurts when you do that, don't do that. MPEG encoders can be configured to only generate I and P frames, thereby reducing the overall delay. Of course, this is a trade-off of compression, and to get the same quality image you need more bandwidth.

There is another MPEG issue. It would be a good idea to maintain audio/video synch, right? Well, the easy way to do this is to delay the audio or video the right amount to keep sync. Depending on the skill of the designer, this too can be a source of delay.

Systems Delay

Beyond the actual MPEG compression and decompression, there is the question about how a product puts the digital video into packets, and the overall buffer management. You can't just dump video into a bucket and send it out on the network only when the bucket is full... that would make you a poor network citizen. It is important to minimize the delay within the protocol handling portions of the system, and to rapidly deal with error conditions introduced by the network.

VBrick Delay

The VB6000 (hardware version -0001 with software release 2.0 and above) has very low delay...about 150 ms.

While there are products that may claim close to 0 delay, they may not be very usable in these modes because they become rather fragile. Frankly, humans cannot really see the difference within about 100 ms except in side-by-side comparisons (which does not happen in real use).

VBrick video delay is such that the VB6000 is excellent for interactive video conferences and for remote control of cameras.

The delay is minimized by creating a GOP using only I and P frames, and by careful engineering of the buffer management system, protocol handling, and through careful adjustment of MPEG timestamps to ensure rapid decoding.

For MPEG-2 the delay would be as low as 180 ms. For MPEG-4, the delay would be as low as 150 ms. For WM, the delay to an existing WM player is about 5 seconds, due to the player.

One last item: StreamPlayer. The desktop MPEG viewer was designed to buffer incoming streams which increases the delay. You will find that StreamPlayer can deliver well under 1 second of delay, depending on the video type and settings.



About VBrick Systems, Inc.

VBrick is the leader in Enterprise IP Video solutions, with over 6,000 corporate, education and government customers and 60,000 installations worldwide. VBrick solutions work over standard IP networks and the Internet to deliver rich media communications that connect people everywhere – from employees and customers, to partners and shareholders. Our comprehensive product suite and end-to-end solutions are used in a wide range of live and on-demand applications including meeting and event broadcasts, distance learning, digital signage, TV distribution, video surveillance, and Web-based marketing campaigns. Headquartered in Wallingford, CT, VBrick's products and services are available through industry-leading value-added resellers.

For more information, visit www.vbrick.com