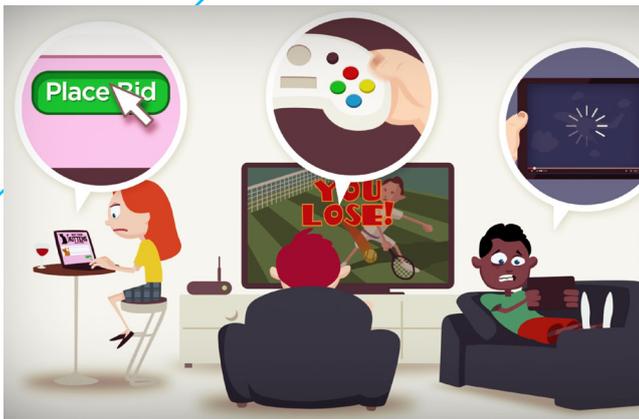# R TE

Latency Fact Sheet

What is latency?

Latency is waiting. Computers and networks does not work infinitely fast. The time it takes before a task is completed is called latency or delay. One such task can be to move data from one place to another over a network. The topic of this factsheet is latency in the Internet.

Some latency cannot be avoided. In a network, data can never travel faster than the speed of light. For instance, a satellite connection has substantial latency, partly due to the distance the signals must cover. We have to live with the speed of light component of latency, but we are   working on minimizing latency caused by other factors.

The effects of high latency

High latency is a problem you may notice when you're using a program like Skype. If someone else in the house is downloading a video at the same time, your Skype connection becomes jerky and may crash. Such extra traffic will also affect gamers who want to play online in real time and companies doing stock trading - both end up buying special and expensive Internet connections to make these work, but often it's not more bandwidth that's needed to improve the speed.
Waiting for webpages that never finishes loading, experiencing the "wheel of death" on YouTube or trying to remote-control a machine where each click takes seconds to register can also be stressful.

Does latency affect business revenues?

Yes! In 2008 the Aberdeen Group found that a 1-second delay in page load time costs 11% fewer page views, a 16% decrease in customer satisfaction, and a 7% decrease in searches that reult in new customers.

In 2009 Microsoft Bing found that introducing 500ms extra delay translated to 1.2% less advertising revenue. Google experimented with injecting 400ms delay before they returned their search page and, based on the drop in new searches, concluded they would lose 0.76% of their revenue by making the experiment permanent.

Can we reduce Internet Latency?

In RITE, we research ways of gradually changing the Internet so that more of the unnecessary delay we experience is removed. This is done by making small, but important, changes to the way computers and networks process Internet data.

The vision of the RITE project is that these new methods that reduces latency will actually be adopted by the industry. Hence, RITE focusses on changing some of the current standards for the Internet. Such changes isn't something that companies like Microsoft and Apple will undertake lightly, but they'll take it up because this has benefits for everyone. If the rest of the world stands up and says   -'ok I get it, we'll standardise it', then RITE has succeeded.
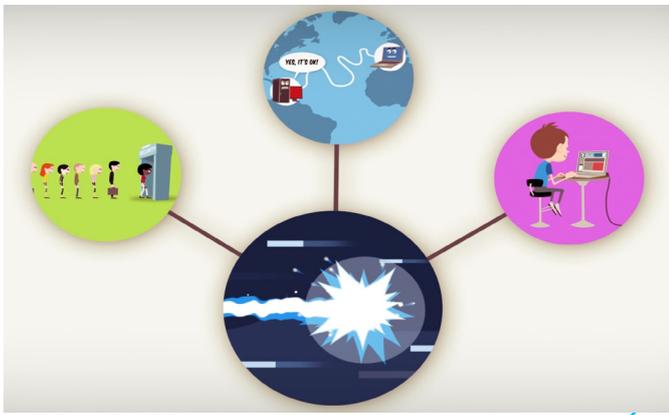
What will lower latency bring?

• Better gaming experience.
• More responsive web applications.
• Better audio and video conference performance.
•Improved interactions with remote-controlled devices
• Better satisfied business customers

What can YOU do about latency?

1. Try to figure out if you get long queues in your home gateway (see activity sheet for how). Contact your ISP to ask for a router that keeps queues short.

2. If available, use a home gateway that supports Active Queue Management.

3. If neither 1 nor 2 is possible, be conscious of what traffic your interactive traffic shares the capacity with. If one family member syncs her Dropbox while you're playing a game, there's bound to be trouble.

Fun fact

Lightspeed in a fibre is only 69% of the actual speed of light between 2 points in vacuum. This is mainly caused by the light bouncing against the walls of the fibre, not going in a straight line. New optical fibres has been developed that make this effect better, but most network connections still use the old, "slow" version.

For detailed information on sources of Internet latency:
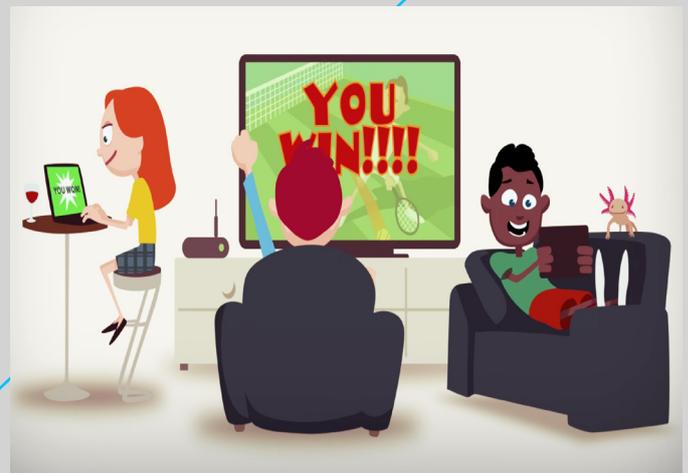http://riteproject.eu/?attachment_id=735

What causes the Internet latency?

Even if all the communication lines you use are upgraded to high bandwidth, the Internet might still seem slow. We point out three reasons for this:

1. In general there will be several handshakes between your device and other computers in the Internet before the real communication can start. For example if you give the name of a computer you want to communicate with (e.g. cnn.com), your computer has to talk to a so called DNS server to find where the computer you want to talk to is located. Then you have to negotiate a connection with the computer you want to talk to. All this adds up to the latency of sending the first byte of real data to or from your computer.

2. Initially the computers talking to each other do not know how fast they can send. They start by sending only a few packets of data (a few thousand bytes). If this succeeds, they gradually send more and more, until their connection is exhausted, and they start to lose data.

3. Since computers in this way usually try to send as much as they can, any bandwidth will be used up until the buffers on the way are completely filled and data is lost. If a computer using this congested path to send an urgent packet (eg. in a voice-over-IP connection) this packet will have to wait in all the congested queues, and will probably arrive too late to be of any value. This situation is called Bufferbloat. In 2010, Jim Gettys of Bell Labs initiated activities under the banner of the term bufferbloat; an allusion to the unnecessary buffer delay that has built up in the Internet.