

End System Traffic Shaping

Traffic Shaping

TCP by its nature creates bursty traffic. Traffic bursts are susceptible to packet loss as there is greater risk of them overfilling a network node's buffer (especially a so called "workgroup switch", namely a relatively cheap switch on the very edge of a network which does not have much memory for each of its interfaces). Such traffic bursts can be avoided if the sender does not send frames 'back to back', but rather adds a short pause in between outgoing frames. This traffic shaping will reduce the maximum transmit rate of the sender (obviously) but by avoiding packet loss will lead to a greater TCP goodput.

Precise Software Pacer (PSPacer)

An example of such a traffic shaper is PSPacer from the [GridMPI project](#). PSPacer is a loadable Linux module and as such does not require any kernel recompilation. PSPacer injects Ethernet PAUSE frames between outgoing frames in order to space them out - this is not the intended use of such frames and as such if PSPacer is used Ethernet flow control should be disabled on the local switch.

-- Main.TobyRodwell - 29 Jan 2007