

PtsCase31

Throughput between JIVE and Shanghai Observatory

PTS Case 31

Synopsis

Project ORIENT has leased an STM-16 circuit between the GEANT2 DK !PoP and the TEIN2 BJ !PoP (in fact, this was an upgrade to a previous STM-4 circuit). The circuit was leased from Telia, but the largest part (Mosco to the Chinese border) is provided by Trans-telecom.

The new ORIENT circuit should have provided the best possible connectivity between JIVE and the Sheshan telescope at Shanghai Observatory (SHAO), but instead it was found that the path was significantly poorer than the GLORIAD path SHAO had been using to connect to JIVE (GLORIAD is routed via the US, and therefore has a much higher latency than ORIENT) - where GLORIAD could achieve a steady 400Mbps data rate, with ORIENT the figure was about 40Mbps (note, the actual eVLBI application rate was 128Mbps, which was in accordance with the observation that eVLBI rate is approx half of the maximum data throughput seen with test applications such as iperf).

eVLBI data is very tolerant to random packet loss, but not very tolerant to delayed data, and so the solution to this problem was to use UDP for the data transfer, instead of TCP. In fact there was nothing new about this case that led to this conclusion, but rather it had just been taking time to enhance the proprietary eVLBI application for use with UDP. Fortunately this upgrade was ready in time for the important Xi'AN APAN Meeting, where delegates saw eVLBI in action and, and the Sheshan telescope was able to successfully send data at a rate of 256Mbps. Of interest, one of the other long-distance eVLBI paths (Australia to JIVE) used a special implementation of TCP, called 'Circuit TCP' (CTCP), which was also successful. CTCP uses the flow control features of TCP, but not congestion control, and is thus ideal for hosts linked with non-contended, point-to-point circuits.

– Main.TobyRodwell - 20 Aug 2007