QoS in the Future Internet

Sara Oueslati and Jim Roberts
France Telecom, DRD/CORE/CPN, Issy les Moulineaux, France
{sara.oueslati;james.roberts}@orange-ft.com

There is increasing pressure on network operators to introduce QoS mechanisms capable of ensuring adequate performance for services like fixed and mobile telephony and IPTV in a ubiquitous converged Internet. Simple reliance on over provisioning is also being called into question as enterprise and residential access rates increase, removing a natural limit on the disruptive potential of individual user traffic. In this presentation we discuss why the move away from the current best effort architecture is proving so difficult, identifying both economic and technical reasons that hinder deployment of QoS solutions.

The present Internet architecture offers little economic incentive for providers to accept the significant additional operational expenses necessary to implement the QoS mechanisms already included in their routers: there is little incentive for a subnetwork to offer better service than its competitors since users are generally unaware of the source of any degradation; users cannot in any case choose an alternative path to the one imposed by QoS-oblivious routing protocols; it proves difficult to achieve the coordination necessary between providers to realize end-to-end QoS guarantees.

Practical inhibitions include the sheer complexity of implementing the QoS models. Since over provisioning appears to meet most user requirements there is understandable reluctance to move away from the operational simplicity of this solution. There is also considerable doubt that the mechanisms of the available Intserv and Diffserv models are actually capable of meeting user requirements if they were deployed.

Through our analysis of the way performance depends on offered demand, the amount of available capacity and the way it is shared, we have been led to conclude that, while it is relatively easy to realize low latency for streaming traffic and high throughput for data traffic, it is impractical to aim for any finer QoS differentiation. Meaningful discrimination can be realized more readily in terms of service availability through selective admission control, as necessary in periods of congestion.

The future Internet might be able to provide adequate QoS through enhancements to the current TCP/IP model, applying the results of work on end-to-end congestion control by Kelly and others. It would be necessary, however, to introduce additional mechanisms to make performance less dependent on user cooperation and to mitigate the impact of overloads. The flow-aware networking approach we have proposed in the past addresses the latter concerns: per-flow fair queuing avoids reliance on correctly implemented end-to-end congestion control while selective flow-level admission control provides the necessary service protection in overload.

Adaptive traffic dependent routing would be a valuable enhancement to the Internet, making it both more efficient and more robust. One recently advocated possibility is to employ enhanced end-to-end multipath congestion control which has been shown to adaptively route traffic over the best paths. This solution remains vulnerable to user misbehaviour and includes no protection against generalized overloads, however. We therefore prefer to advocate a flow-aware solution where fair queuing and admission control are used to control both bandwidth sharing and path choice.