Terms of Reference for a Quality of Service Think Tank

UKERNA
November 2000

Contents

1 Introduction

The challenge of large-scale network quality of service support on SuperJANET4 is formidable. The first step towards tackling this challenge is to harness the expertise available within the academic community and from the supplier of the switching equipment used within SuperJANET4 to create a roadmap for the development of network QoS. To this end, this document specifies the Terms of Reference for a network QoS Think Tank whose job will be to produce the development roadmap.

2 Background

The following trends are apparent on today's Internet:

- Many new Internet-enabled applications are emerging, for example pervasive videoconferencing and video streaming services, virtual networked environments etc.
- Convergence of other networks such as telephone, radio, television with the Internet is underway.
- Network traffic is increasing as the number of Internet users and applications increases. Capacity upgrade needs to be complimented by introduction of complimentary network traffic engineering techniques, to improve utilization, and to tailor services to user needs.
- There is increasing diversity among user application needs, and the network needs to be engineered beyond a single level of service in order to meet this greater range of requirements.

In order to keep pace with these developments increasing Internet bandwidth on its own is not going to be sufficient, but clearly it must happen to avoid traffic congestion. This effectively means making the transition from single service and best effort service networks that exist now into multi-service networks which provides different levels of service.

IP design inherently provides a ‘best effort’ service which is subject to unpredictable delays and data loss. As stated above, congestion can be avoided by increasing bandwidth, however many new Internet applications are multimedia and not only require a lot of bandwidth, but also have strict timing requirements requiring more than simple best efforts.

Traffic Engineering in this context is a term used to describe how different types of traffic are managed over an IP network. Traffic engineering allows varying levels of network resource to be allocated to different applications based on their performance requirements or by customer allocation i.e. they have paid for a high priority service. There are two fundamental elements of traffic management:

1. The Quality of Service (QoS) algorithms, protocols and policies that enable differentiation of IP packets and the application of appropriate behaviour on a particular traffic flow. For example techniques exist in IP routing technology to mark packets; place packets into queues that are given weighted priority; signal the network to reserve network capacity for end-to-end QoS; match traffic to available network capacity (traffic shaping); and perform network congestion anticipation.

2. Use of service engineering techniques such as caching, multicasting, Multiprotocol Label Switching (MPLS) to reduce, redirect and balance network traffic.

Ultimately traffic engineering should provide a reliable end-to-end Quality of Service to the user. End-to-end however means providing service delivery guarantees that span multiple network management domains and in the context of SuperJANET4 the management domains are campus LANs, MANs, SuperJANET backbone. In terms of implementing end-to-end QoS, it not only means that standard QoS technologies must be deployed, but that service level agreements and operational standards must be in place across the multiple management domains.

3 QoS and SuperJANET

It is important that SuperJANET provides a platform on which a range of existing and future network applications can run to specified service levels and to do this the following areas must be developed and implemented:

- Implementation of QoS technologies to control traffic behaviour
- Manage traffic effectively using techniques such as caching and multicasting
- Implement a policy framework which the QoS technologies carry out, allowing an end-to-end predictable service

Each of these areas must be implemented to current standards and track emerging standards to assist peering across multiple network domains.

4 Think Tank Deliverables

In order to define the foundations on which each of these areas can be developed, a think-tank focusing on requirement for QoS within the JANET network is proposed. The requirement for this think tank is driven as much by a perceived user need for end-to-end (host to host or host to server) QoS as by the increasing demands of existing and developing network applications.

There are three deliverables required from the think-tank and these are:

1. Assess the requirement for QoS on JANET
2. Develop a policy framework on which to implement QoS on JANET
3. Provide initial recommendations for the technical mechanisms by which to implement QoS on JANET

As these deliverables are developed a process of wider consultation will need to take place. This consultation will be with representative experts from the following areas:

- End users from a broad cross-section of the JANET community.
- MAN and institution service managers.
- Experts in technology and JANET policy.

5 Think Tank Membership

The think tank will have a membership of:

- Representatives from the sector actively involved in QoS research (Chris Cooper, RAL, Tim Chown, University of Southampton and Chris Edwards, University of Lancaster)
- A representative from an application that has QoS requirements (e.g. videoconferencing) (David Price, University of Wales, Aberystwyth)
- A representative from the JANET Network Operations Centre (NOSC) (Jonathan Couzens)
- A representative from the research GRID community (Robin Tasker, Daresbury Labs)

- A representative from the of the switching equipment used within SuperJANET4 (Cisco Systems Ltd) (Jane Butler)
- Representatives from UKERNA (Jeremy Sharp and Rina Samani)

The group will be supported by a UKERNA project manager who will be responsible for producing the three deliverables.

6 Timescales

It is expected that the think-tank will exist for a period of not more than 4 months and in that time it will meet three times: Once to brainstorm each area of work; a second time to review progress and deliverables; and a third time to approve the deliverables. It is expected that the group will report by the end of April 2001. The meeting timings are as follows:

- Meeting 1: December 2000 / January 2001
- Meeting 2: End February 2001
- Meeting 3: Late April 2001