QoS in Hybrid Networks – An Operator's Perspective

Aiko Pras[1], Remco van de Meent[1], and Michel Mandjes[2]

[1] University of Twente, The Netherlands
{pras, meent}@cs.utwente.nl
[2] Center for Mathematics and Computer Science, The Netherlands
michel.mandjes@cwi.nl

Abstract. The goal of this paper is to foster discussions on future directions for QoS related research. The paper takes the viewpoint of an operator; as an example it presents the topology, capacity and expected usage of the next generation research network within the Netherlands, called SURFnet6. In that network the traffic from normal and demanding users gets separated; the mechanism to realize this is lambda switching. The key method to assure performance is overprovisioning; there is no need to use DiffServ or IntServ.

1 Introduction

Since many years Quality of Service (QoS) is an important research topic. In literature, QoS has been defined as "providing service differentiation and performance assurance for Internet applications" [1]. Traditionally, QoS research relied on frameworks like IntServ and DiffServ to achieve its goals. The Internet hype, which led to a gross overinvestment in transmission capacity, and the subsequent collapse of the bubble, require us to rethink these frameworks, however. In the last five years we've witnessed a growth of available backbone capacity that exceeded the growth of Internet usage, which remained stable at approximately 100% per year [2]. As a result, prices went down considerably; for the costs of a few kilometres highway, it is now possible to create a nation-wide backbone in the Tbps range. Backbone link capacity need therefore no longer be regarded as a scarce resource.

Section 2 of this paper discusses the consequences of these developments for the next generation of the Dutch research network, called SURFnet6. It shows how service differentiation and performance assurance will be realized in that network. Section 3 provides the conclusions and identifies some remaining QoS related research challenges.

2 SURFnet6

The organization responsible for the Dutch research network is SURFnet. The first generation of their network was installed in the early eighties of the previous century and installation of the sixth generation started at the end of 2004. The main idea

H. de Meer and N. Bhatti (Eds.): IWQoS 2005, LNCS 3552, pp. 388 – 391, 2005.
© IFIP International Federation for Information Processing 2005

behind this new network is to take advantage of the dark-fibre infrastructure that has recently become available, and to use lambda switching to separate the traffic generated by demanding users from that of normal users. In the next subsections the topology, capacity and expected usage of that network will be discussed.

2.1 Topology

It is exactly five years back that the Internet hype reached its top. At that time it was impossible to rent dark-fibre. Since there was a common believe that Internet usage would explode, many organizations started to install their own fibre infrastructure. In the year 2000 the cost of digging fibre trenches was in the range of 50 to 100 euro / meter. Soon after the collapse of the Internet bubble, it became clear that there had been an over-investment in fibre capacity. As a result of this over-investment, a significant part of the available fibre capacity remained unused. The owners of the fibre infra-structure were subsequently forced to change their business models and, as a result, it is now possible to rent dark-fibre from multiple parties. The prices for dark-fibre are, in rural areas, in the range of a few euros per meter (for a 15 years lease period). Also digging costs went considerably down to around one-fifth of the original costs. These developments made it possible for SURFnet to acquire a dedicated fibre infrastructure for the new SURFnet6 research network (Figure 1). The length of this infrastructure is around 5300 km, which is comparable to the size of the Dutch rail and highway infrastructure.

Fig. 1. Topology

2.2 Capacity

With Coarse Wave Division Multiplexing (CWDM), which is a relative cheap technology, it is possible to use 16 wavelengths on a single fibre pair. If OC-48 is used, the resulting capacity becomes 40 Gbps. Besides some CWDM rings in metro areas, SURFnet will use the more advanced Dense Wave Division Multiplexing (DWDM) on five individual rings of dark fibre pairs, with a capacity up to 720 Gbps per fibre pair per ring. Note that some links within SURFnet6 will not consist of a single fibre pair, but of multiple pairs. Hundreds of such pairs may be included within a single duct.

2.3 Usage

The usage of SURFnet6 is expected to keep on doubling every year [3]. Two kinds of users can be distinguished:

- Normal users, who use the network for web surfing and email exchange. Some of these users will also participate in P2P networks and watch Internet videos.
- Demanding users, who would like to exchange data at Gigabit speed for longer periods of time. Examples of such users are physicists, who perform nuclear experiments, and astronomers, who correlate data from different radio telescopes (like LOFAR).

To assure performance, no special measures are required for the normal users; there is no need for relatively complex mechanisms like DiffServ and IntServ. As explained in the previous subsection, the issue is not a lack of link capacity between network nodes. Rather the issue is the complexity of the nodes themselves, which should be kept at a minimum. Contrary to what some researchers claim, overprovisioning may be a viable method to assure performance!

For demanding users the situation is somewhat different; the amount of data that these users exchange may be such that, without special measures, performance for normal users can no longer be guaranteed. For that reason traffic from the relatively small number of demanding users will be separated from that of normal users. The mechanism to realize this separation is lambda switching. Dedicated optical path connections are established between these demanding users, using dedicated DWDM colours. This separation has the additional advantage that demanding users are free to operate dedicated high-speed transport protocols; such protocols could disturb the correct operation of the 'normal' TCP protocol and should therefore not be mixed with normal TCP traffic on top of a single IP network. It is important to note that optical path establishment is not limited to the Netherlands, but via NetherLight [5] and the Global Lambda Integrated Facility (GLIF) [6] it is already possible to reach major parts of Europe, the US, Canada, the Far East as well as Australia.

3 Conclusions and Remaining QoS Challenges

Since link capacity will continue to grow faster than the capacity of routers, SURFnet doesn't see a need for IntServ and DiffServ. Our belief is that the shift towards optical networks should also lead to a shift in thinking about QoS: instead of ensuring QoS at a fine level of granularity (packets and flows), QoS should be ensured at an aggregate level (provisioning of links and optical paths). In [4], for example, a study is described that balances overprovisioning and resource demands by looking at peak usage at (small) time scales that correspond to users' perception of QoS. The results may be used to determine whether investments in link upgrades are warranted or may be postponed without sacrificing performance. Although the study focuses on access links, the ideas could also be extended to a network wide setting. Also automatic lambda management may be an interesting topic for research: when is it worthwhile to setup and use a lightpath between two points in the network, and how can this be done automatically?

Acknowledgement

We would like to thank Erik-Jan Bos, Jan van Oorschot (both SURFnet) and Remco van Mook (Virtu Secure Webservices) for their input, comments and valuable information.

References

1. W. Zhao, D. Olshefski and H. Schulzrinne: Internet quality of service: An overview, Technical Report CUCS-003-00, Columbia Univ., Computer Science Dept., Feb. 2000
2. A. M. Odlyzko: Internet traffic growth: Sources and implications, in: Optical Transmission Systems and Equipment for WDM Networking II, Proc. SPIE, vol. 5247, 2003
3. SURFnet Annual report, 2003, http://www.surfnet.nl/staging/attachment.db?349156
4. R. v.d. Meent, A. Pras, M.R.H. Mandjes, J.L. van den Berg, F. Roijers, L.J.M. Nieuwenhuis, P.H.A. Venemans: Burstiness predictions based on rough network traffic measurements, Proceedings of the 19th World Telecommunications Congress (WTC/ISS 2004), Seoul, Korea, September 2004
5. NetherLight home page, http://www.netherlight.net/
6. GLIF home page, http://www.glif.is