

QoS deployment in Regional Networks

In this section we address the deployment of QoS in the Regional Network based on the experience gained by QoS operation and trialling in the project. Due to the range and scope of network infrastructures that operate in this area, it is difficult to present a single case for QoS that fits all cases. Instead, this section will present a discussion of the issues that may be encountered when considering QoS in the Regional Networks and how this can interact with the JANET QoS Policy presented above. In particular, we will address this based on the two broad types of connectivity that are offered to sites from Regional Networks; High Bandwidth Access, such as that found in universities, and Low Bandwidth Access for schools and colleges. The former will include issues related to QoS in the network core while in the later we address QoS issues at the Edge.

This section will first discuss the architecture of the network, the QoS-enabled services that will be supported and the technologies employed in the network. Thereafter, we will examine the operational issues of QoS such as the policy employed and the monitoring mechanism used. We will then conclude with a discussion of over-provisioning in the Regional Network as an approach to QoS and Managed Bandwidth Services.

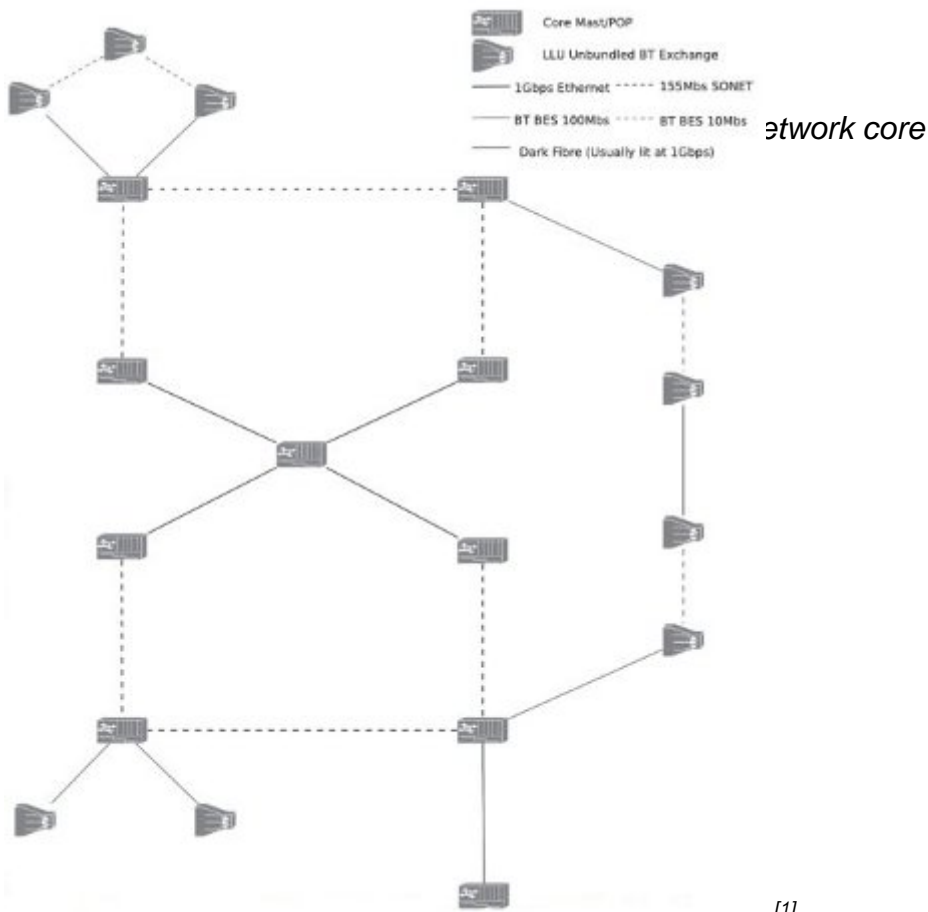
4.1 Regional Network Architectures

The differing scope and range of coverage of Regional Networks in the UK makes it impossible to identify one general case architecture that can apply to all RNOs. In all cases, however, we assume there to be a high-capacity core that connects large end sites and a collection of more heterogeneous edge networks offering connectivity to smaller sites. We discuss this in more detail here to provide a flavour of the composition of each aspect of the network.

4.1.1 Network Core and High Bandwidth Access Sites

Network connectivity in the core (or backbone) and that provided to large university sites will typically be in the region from +100Mbit/s up to 1Gbit/s and beyond. It is also not unusual to see multiple high capacity Points of Presence with multiple levels of redundancy enforced between them. As discussed in section 3, we could also assume, with some validity, that this aspect of the Regional Network will be generally overprovisioned for the purposes of QoS and therefore that explicit support is unnecessary in some cases.

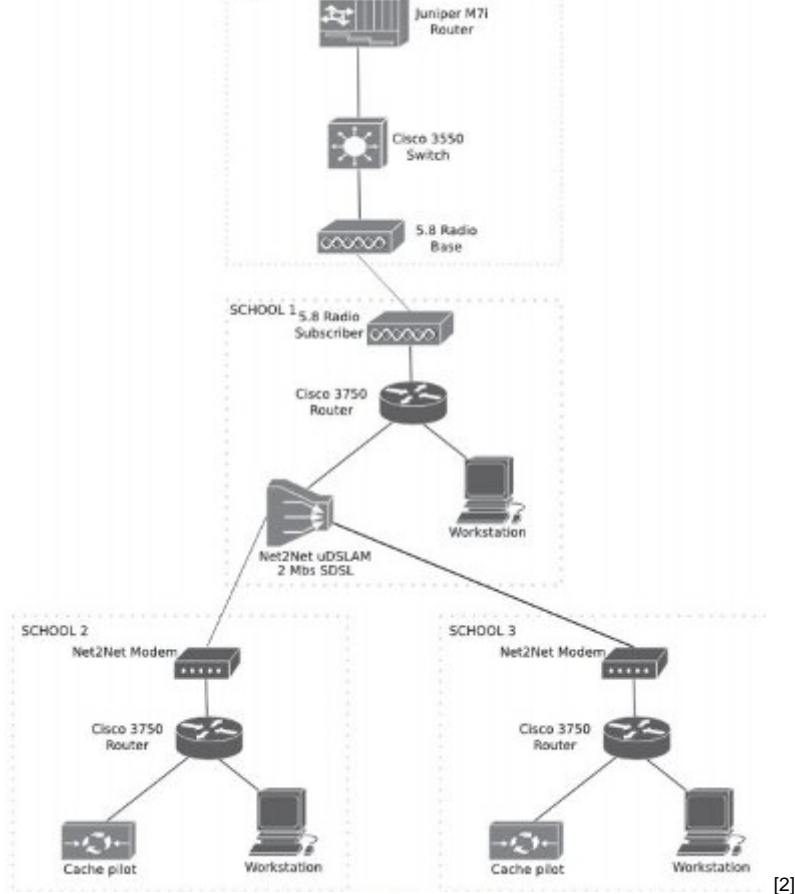
Since trialling of QoS deployment in this area was conducted as part of the QoS Development Project, we will use this section to present an overview of our findings for explicit QoS usage in this area. In particular, we will highlight equipment configurations where possible and discuss traffic marking schemes employed by the partners. Figure 4-1 presents a summary of the Regional Network core employed by Cumbria and Lancashire Education Online (CLEO) at Lancaster [LANCP2].



4.1.2 Network Edge and Low Bandwidth Access Sites

The other aspect of supporting QoS in Regional Networks is concerned with connecting sites on low bandwidth connections, typically schools and colleges. In this context, we define 'low bandwidth' as being anything under the threshold for high bandwidth (100Mbit/s) but a more common connection speed would be in the region of 10Mbit/s or lower. In contrast to the larger sites, schools may be connected at the Regional Network edge via a range of 'last mile' technologies such as ISDN, DSL, fibre or radio. Figure 4-2 gives an example of how schools might be connected, again based on the CLEO network.

Figure 4-2 - Smaller sites connected at the Regional Network edge



QoS mechanisms are generally not seen as an answer to low bandwidth problems on the network edge as a school with only a 128Kbit/s ISDN connection is going to have problems with videoconferencing and rich media streaming no matter how much traffic shaping is applied. Bandwidth issues typically arise in Local Authority (i.e. edge) networks due to low bandwidth ISDN and ADSL connections. For example, it is estimated that as much as 30% of all schools in Scotland have links with 2Mbit/s or less bandwidth. However, where a site has sufficient bandwidth (typically 4Mbit/s or better), QoS techniques can be used to support VoIP and other traffic types.

4.2 Resource Provisioning

When considering QoS in the Regional Network, it is important to identify and classify the types of traffic that will be seen on the network to determine how it will be handled on QoS-enabled links. This section first classifies the traffic on Regional Networks, discussing the priority that should be afforded to each, before attempting to aggregate this into a number of more simple groups as was done in section 3 to enhance the scalability and compatibility of the approach.

4.2.1 QoS Traffic Classification

As in most networks, a number of broad traffic types can be identified based on the typical load seen on the Regional Network:

- **Data transfer** - Besides normal web traffic, end sites frequently need to upload and

exchange content with servers located elsewhere in order to backup data to an external source or download new content. This is not a critical operation, nor is it sensitive to network conditions, and thus can be given a low priority.

- **Web traffic** - This class represents normal web requests from machines connected to the network. We can expect this class of traffic to represent the majority of the background load on the system during normal working hours. This traffic is fairly important but can still be classified as normal priority.
- **VoIP traffic** - VoIP is increasingly being used both on a personal level and more formally as a research tool in end sites. Audio traffic typically requires around 64Kbit/s and is sensitive to network conditions. As such, VoIP traffic should be given a higher priority than normal traffic where possible.
- **Videoconferencing traffic** - Increasingly, both large and small sites may wish to hold videoconferencing calls with other sites regionally, nationally and beyond. A videoconferencing call typically needs around 320Kbit/s in both directions for video traffic and 64Kbit/s in both directions for the audio traffic. This traffic needs to be highly prioritized due to the strict bandwidth requirements and its delay sensitive nature.
- **Network Control traffic** - Routing information and other network control traffic should be classified as high priority as it is used to exchange state information between routers and other network devices. However, this usually utilizes only a limited amount of bandwidth and so introduces only a minor overhead.

Other Traffic Types

Cache Pilot Updates - As outlined in the Lancaster deliverable to the JANET QoS development project [LANCP2], some schools are allocated a piece of equipment called a cache pilot. The cache pilot is configured to download and cache content for websites and resources that the school accesses frequently as a proxy. Cache pilots are typically configured to download their content during the night in order to minimise their impact during the working day. The updates to the cache pilots, while usually voluminous, are not critical and can therefore be given low priority.

Multicast applications and services - While IP multicast has well known scalability issues when applied on a large scale, in the limited scope of a Regional Network there is great potential for this technique to supplement QoS and reduce bandwidth consumption. There are a number of applications and services that are commonly associated with multicast in this scope. These include:

- **Real-time video distribution.** The real-time or live nature of the event means all viewers have synchronous requirements for data delivery. A major example of this is (commercial) IPTV type services.
- **Real-time audio distribution.** As per video, but audio only (radio stations, conference audio, for example).
- **Real-time conferencing applications.** In this case multiple participants may use multicast for live conferencing.

Since the multicast services listed above typically encompass real-time, delay sensitive traffic, multicast applications in this case should be classified as higher priority.

4.2.2 Traffic Class Aggregation

Based on the above classifications, we can group traffic types based on the service they expect to receive in the network in order to arrive at a small number of classes that is more straightforward to enforce and scalable to implement. As a general rule, we can identify 3 classes for high priority traffic (Premium IP), normal traffic (Best Effort) and low priority traffic (Less than Best Effort). As such, the traffic types from section 4.2.1 could, for example, be grouped as follows:

- **Premium IP** - VoIP, videoconferencing, multicast services (as discussed in section 4.2.1), network control
- **Best Effort** - web
- **Less than Best Effort** - data transfer, cache pilot updates

This process is also useful to highlight that, unless applications with specific QoS requirements are being considered as part of the QoS Policy, there may not even be a need for Premium IP traffic marking in the Regional Network. This was found to be the case in the Kentish MAN where an investigation was carried out on the network to establish the different types of traffic. No premium traffic was found but all connected sites were actively consulted on the classification of their traffic and the consensus was that there should be no traffic classed as premium. Subsequently, all traffic was remarked to 'Best Effort' (BE) with the exception of the Mirror Service which was remarked as 'Less than Best Effort' (LBE). Thus, in the event of congestion, normal traffic would continue to be serviced but Mirror Service traffic would only be guaranteed a small percentage necessary to ensure TCP sessions remained established.

4.3 Regional Network QoS Implementation

As part of the QoS development project, a number of Regional Network Operators experimented with implementing QoS on their infrastructure and gained useful experience on the practical issues involved in deploying and operating QoS on their equipment. This section will discuss these experiences in terms of deploying QoS on the core and edge equipment.

4.3.1 Regional Network Core Equipment

This aspect of the Regional Network will be similar to many ways to the JANET network as discussed in the previous chapter. The network infrastructure is likely to compose of high capacity packet-switched links in the region of 1Gbit/s - 10Gbit/s interconnected by core routers. During the JANET QoS development project, a number of Regional Network Operators experimented with QoS provisioning in this scope and the results can be found in the appropriate partner deliverables. In particular, the deliverables from Kentish MAN [KentP2], Lancaster [LANCP2] and SSDN [SSDNP2] will be of interest here. These deliverables also include some practical examples of how QoS can be configured on commonly used router equipment used by the partners including the Juniper® M7i, Cisco® 7401 and Cisco® 3550. We have also included these configurations for convenience in Appendix A2, A3 and A4.

QoS deployment in this scope would be based on the DiffServ model with several priority

queues implemented on network routers and packets belonging to higher-priority traffic classes with elevated QoS. As discussed in section 3, this would most likely involve static provisioning based on explicit user requests and would be policed with ingress filtering at the network border routers based on traffic class and source/destination IP addresses.

Based on the JANET QoS Policy [JANETPolicy] and the findings of our QoS trialling, the favoured approach to QoS in the Regional Network core seems to be based on overprovisioning except in specific circumstances. If so, the network core can again be regarded as QoS-neutral in that traffic marking is not explicitly honoured but neither is it dropped or remarked. Exceptions to this case may be required if, for example, the Regional Network Operator receives an explicit request from a site to support QoS-enabled applications such as Janet Videoconferencing. In this case, a site may require a specific bandwidth reservation to be made for this which the Regional Network should support. If so, the Regional Network border routers can again be configured to provide admission control based on the packet source/destination addresses and traffic marking.

4.3.2 Access Technology Considerations

In contrast to the network core (and high bandwidth sites) which are typically connected via fibre, low bandwidth sites can be connected at the network edge using a range of technologies including DSL, ISDN and radio in addition to fibre. A brief overview of these access technologies is given below.

ISDN

ISDN (Integrated Services Digital Network) is a circuit-switched telephone network system, designed to allow digital transmission of voice and data over standard copper telephone wires. It offers circuit-switched connections (for either voice or data) in increments of 64Kbit/s up to 128Kbit/s and was common before DSL became the prevalent broadband technology. ISDN is largely considered to be a legacy technology now but low bandwidth sites may still be connected to the Regional Network edge via ISDN.

DSL

DSL technology comes in many forms and can be used to provide a broadband connection over standard copper telephone lines. Download speeds vary between 144Kbit/s to 24Mbit/s depending on the DSL technology used and the length and quality of the copper line. The distance up to which most DSL technologies work is 3km, though it can be extended, albeit at lower speeds.

- **SDSL** - Symmetric Digital Subscriber Line is the version of DSL where both the upload and download speeds are the same. This can vary from 144Kbit/s to 2320Kbit/s. SDSL cannot be shared with a phone line and requires a dedicated copper line. However due to the symmetric bandwidths, this enables schools to upload content to servers at a higher bandwidth.
- **ADSL** - Sites connected via ADSL2+ have a typical upload speed of 1Mbit/s and a download speed of 2.3Mbit/s. However if sites reside on telephone lines longer than 3km, they might only be connectable at lower speeds.

Radio

Radio links can be used to connect schools via masts or 'chained' together with other schools to reach isolated areas. Radios are highly useful to connect schools where a fibre run or DSL is prohibitively expensive (e.g. in hilly Cumbria) and where the site has line of sight to a mast or another school. One example of where this has been used effectively is in the CLEO (Cumbria and Lancashire Education Online) network that covers an extensive, geographically 'difficult' area. CLEO operates equipment in the 2.4GHz and, more recently, the 5.8GHz radio frequencies. The 2.4GHz radios work up to distances of 15km and have a bandwidth of around 11Mbit/s. However, these radios are being replaced by 5.8GHz radios which have higher bandwidth and fewer interference problems. The 5.8GHz radios also have a bandwidth of around 24Mbit/s and a greater range of around 20km.

Our experiments do show that while QoS provisioning and deployment is possible for edge network technologies, the sheer range of available vendor equipment and possible configurations makes it difficult to guarantee. For example, we found that while many technologies support QoS to some extent, certain devices either exhibit inconsistent behaviour with QoS enabled or fail to work at all. As such, we recommend that a thorough evaluation be conducted on all equipment before QoS deployments are seriously considered. Please see the individual partner deliverables for a more detailed description of QoS operation on the network edge.

4.4 Monitoring and Measurement in the Regional Network

In order to offer a production level QoS-enabled service on the Regional Network, it is necessary to collect network performance metrics to demonstrate that the level of service being provided meets the agreed Service Level Agreement (SLA). At the inception of the QoS development project these metrics were not being collected and so it was necessary to deploy appropriate monitoring devices at physical locations in the network in order to collect and store the necessary network performance information.

4.4.1 Monitoring Granularity

Monitoring functionality can be applied on several levels depending on the level of QoS being deployed. At a high level, the Regional Network Operator will need to employ some form of monitoring simply to maintain control of its network core. This could include functionality to measure the total load throughout the nodes in the core to identify congestion points or determine failures. To support the deployment of QoS, this can be extended to measure the aggregated throughput of the supported traffic classes and ensure that their requirements are met. This should ideally be performed dynamically but may also be achieved by taking snapshots of network load at various points of the network. As discussed in section 3, monitoring functionality has yet to fully embrace QoS parameters and so some level of improvisation may be necessary.

If QoS provisioning is to be performed for specific applications or traffic classes, the Regional Network may extend its monitoring functionality to measure the metrics of this in its network as part of the end-to-end process. In this case, in addition to throughput, the operator may also attempt to monitor specific QoS metrics for this flow such as loss, delay and jitter as

shown in Figure 4-3 overleaf. If over-provisioning is used as the QoS-support approach, monitoring in this case may only extend to monitoring dynamic link-utilisation as in section 2.4.1.

4.4.2 Monitoring Janet Videoconferencing

JANET had previously developed a monitoring system to record network performance and it was decided to deploy this system to support Janet Videoconferencing (see section 3.3.1). The monitoring system is based on a Cisco® 1760 router running the Service Assurance Agent (SAA) module. This is located at the site to be monitored and reports network performance information back to a central server, which stores the information and presents meaningful graphs of network performance.

Monitoring units were preconfigured for installation at the location of each of the Janet Videoconferencing MCUs, namely the JANET co-location facilities at Reading and Leeds, and for the Janet Videoconferencing Management Centre in Edinburgh. Janet Videoconferencing allocated IP addresses within the appropriate subnets and allocated a switch port for each of the monitoring devices. The JANET Network Operations and Service Centre (NOSC) applied appropriate changes to Access Control Lists to permit the monitoring devices to communicate with the back-end of the monitoring system.

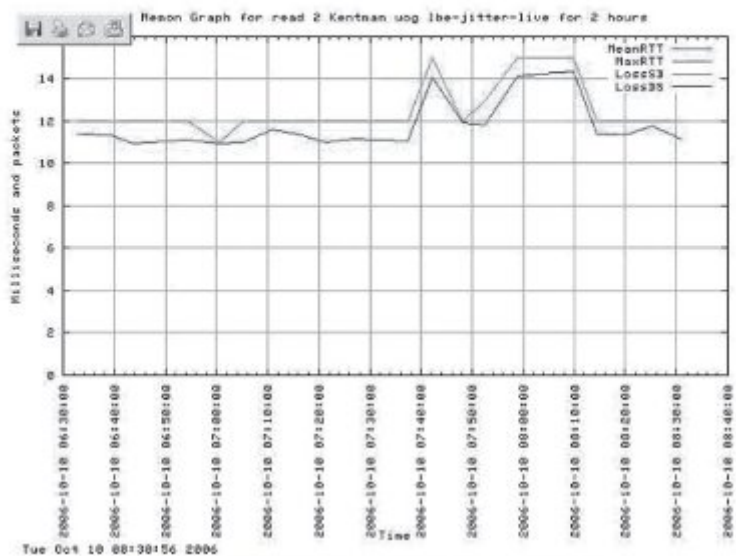
4.4.3 Multicast Monitoring

In the above sections we have discussed measuring unicast QoS using tools that measure various metrics such as packet loss. In the JANET QoS Project, we also considered multicast QoS measurement in this scope. Multicast can be a complex protocol to deploy, manage and monitor, but the benefits of using it can be great. Typical examples of multicast traffic are video or audio streaming, but it can also be used for file distribution, distributed backups or distributed configuration changes. Perhaps the most high profile use of multicast in UK networks is for AccessGrid®, a high quality conferencing system built on the original vic and rat multicast tools.

Figure 4-3 - Packet monitoring graphs for QoS parameters

JANET Measurement

The dataset is: QoS Project
The duration is: 2 hours
The filename is: read_2_Kentman_uog_lbe
Graph type is: Normal

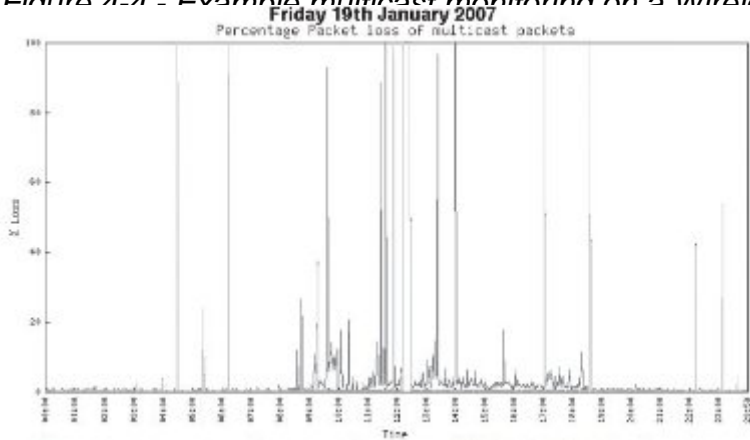


[3]

A number of multicast monitoring tools already exist and are in use including ssm ping [ssmping], mtrace [mtrace] and, for the purposes of QoS multicast, dbeacon [dbeacon]. The dbeacon tool is the common method for monitoring QoS between multiple multicast nodes in a system. With this tool, sites that wish to monitor multicast transmission and reception deploy a number of dbeacon nodes (ideally on the same links as the multicast traffic) which form a monitoring matrix to report on traffic flows in the network. Dbeacon includes support for historical views of connectivity and can show QoS metrics such as delays, loss or jitter.

Within the context of the QoS Project, Southampton did some further research on QoS multicast monitoring and developed a tool called Quasi [Quasi] which can be used to monitor multicast streams dynamically on real-world networks. Further information on the Quasi tool can be found in the Southampton QoS Project deliverable [Quasi] and a typical example of Quasi monitoring output in a wireless network is given in Figure 4-4 below.

Figure 4-4 - Example multicast monitoring on a Wireless LAN at Southampton 4.5.1 QoS



[4]

4.5 Regional Network QoS Policy

In the Regional Network scope, the QoS Policy will primarily define the traffic marking and handling procedure. As with the JANET network core, DiffServ had been adopted as the de facto protocol to use in this area. As we have seen, the JANET approach to QoS is to rely on over-provisioning in their network core to make it QoS-neutral and DSCP transparent. This effectively defines that traffic flowing through the network should not be given priority based on its class but, at the same time, it should not be dropped, reclassified or modified either.

While this approach might ensure that all traffic traversing the network is given sufficient resources to meet its requirements, it has a disadvantage as it breaks the end-to-end QoS provisioning model. This is a critical aspect of using QoS as, without end-to-end service guarantees on traffic handling, the DiffServ approach cannot be fully enforced. In the QoS development project, the Regional Network Operators individually experimented with traffic marking and handling policies in their networks. We use this section to present an overview of the policies adopted in this case.

4.5.1 QoS DSCP Value

QoS traffic is marked in DiffServ by setting the DSCP (or Class of Service) field in the IP packet. Within the IETF and elsewhere, various values have been defined and used based on the service classes to be enforced (EF, AF, BE, BLE) [RFC3260]. The QoS Project did not deem it necessary to formalise the DSCP marking procedure for traffic on JANET but does offer recommendations in a separate document [JANETDSCP].

For testing purposes, and to simplify deployment, the aggregated traffic classes identified in section 4.2.2 were used as the basis for the marking scheme employed by the project partners. This model was enforced by Kentish MAN for example who identified the 3 traffic classes for their trials as follows:

Traffic Class	DSCP Value	Cos Value	Description
Premium	46	5	Strict priority
Best Effort	0	0	Delivery with no
Less than Best Effort	8	1	Moderate bandwidth guarantees

This was also reinforced by the marking scheme adopted by Lancaster which supplemented this with an additional class specifically to identify Network Control traffic:

Traffic Class	DSCP Value	Binary Value
Network Control	34	100010
Premium	26	011010
Best Effort	18	010010
Less than Best Effort	10	001010

The need to simplify the DSCP marking scheme was further highlighted by issues encountered on some equipment while configuring QoS. We found that several implementations do not offer more than 2-3 separate queues for traffic handling which obviously limits the number of unique traffic classes that can be enforced. As such, our experience is that a limited number of classes makes sense, both in terms of simplifying the QoS policy and ensuring it can be enforced practically on the network.

4.5.2 Over-Provisioning

Based on the experience of the individual partners during the QoS project, there is an increasingly compelling case to support over-provisioning as the most practical approach to QoS deployment in the Regional Network. While this does not enforce the end-to-end model of QoS provisioning, the natural over-provisioning that is typical in Regional Network cores combined with the potential issues related to vendor equipment makes it the most appropriate near-term solution for supporting QoS. This approach also reinforces the JANET QoS policy which recommends that while no explicit QoS provisioning is adopted, the network should be made QoS-neutral such that traffic marking is not handled separately but is not dropped or altered either. In this case, the Regional Network border routers can also provide admission control based on the packet source/destination address and traffic marking but no further action is taken thereafter.

4.5.3 Managed Bandwidth Service

As discussed in section 3.4, JANET Lightpath is an optical infrastructure, centrally managed by the JANET NOC infrastructure which spans the JANET core and ends at Regional Networks Entry Points (RNEPs). As such, when a Regional Network Operator

(RNO) needs to extend a Lightpath circuit to an end user site to provide a managed

bandwidth service which co-locates with a RNEP, it is relatively straightforward: it normally simply needs patching between racks (or between buildings). This often occurs as RNEPs tend to be located at universities where users of a bandwidth channel such as researchers are located.

However, a RNO might face problems when it is necessary to extend a bandwidth channel to a user site which is not co-located with the RNEP where the core span of the bandwidth channel ends. Several options may be used in this case and the choice depends on the RNO's capabilities, the channel bandwidth and a number of potential users of such channels.

- **Leasing a circuit of the required bandwidth from a third-party provider.** This is a reasonable choice when the RNO has no infrastructure of its own to provision a channel of required bandwidth and an estimated number of potential users of such channels in the foreseeable future is close to zero.
- **Provision a channel through the RNO's own DWDM/SDH/OTN infrastructure.** This is a good way of extending a core span of the channel for several reasons:
 - in this case the RNO's span of the channel will be built on the same high-speed circuit-switched technology as the core span, providing consistency in the channel quality (deterministic bandwidth), a high degree of protection for production traffic (because of the nature of circuit-switched networks), and high scalability of channels in terms of their speed and numbers
 - the same RNO's Layer 0/1 infrastructure is shared for both IP production and bandwidth channel as in the JANET core case
 - the management of the channel will be simpler than in the case of the leased circuit because it will exclude the interaction with the completely external organisation.

Of course, if the RNO doesn't have such an infrastructure underpinning its IP layer, the research centre's need for a few bandwidth channels can hardly justify its creation.

- **Provision a channel through the IP production layer of the Regional Network.** This option is generally feasible only for medium speed channels as the required amount of bandwidth should not exceed some reasonable percentage of the Regional Network's links capacity. Theoretically any upper limit of reservation up to full link capacity is possible when the IP network is used for establishing bandwidth channels. However, in practice an upper limit of reservation should be chosen such that: 1) research flows have enough space not to starve each other; 2) if they share bandwidth with production traffic then the latter also should have sufficient capacity. Generally, this means that the upper limit of reservation for bandwidth channels should be in the 20% range of a link's capacity. This speculation proceeds from an assumption that IP links are utilised by the production traffic no more than 30% so that an additional 20% utilisation by research traffic will not raise utilisation above 50%, which is normally considered as an upper limit of average utilisation (see section 2.4 for details).

In practice this means that only a couple of 1Gbit/s bandwidth channels will be allowed on 10Gbit/s interfaces or a couple of 100Mbit/s bandwidth channels will be allowed on 1Gbit/s interfaces. If the RNO decides to implement a bandwidth channel span through its production IP network and the core of this network is over-provisioned, the policy and recommendations given in sections 3.1 and 3.2 might be of use.

Source URL: <https://community-stg.jisc.ac.uk/library/janet-services-documentation/qos-deployment-regional-networks>

Links

[1] <http://community.ja.net/system/files/images/tg-qos-4.1.jpg>

[2] <http://community.ja.net/system/files/images/tg-qos-4.2.jpg>

[3] <http://community.ja.net/system/files/images/tg-qos-4.3.jpg>

[4] <http://community.ja.net/system/files/images/tg-qos-4.4.jpg>