

JANET QoS Development Project



JISC

**Investigating the Development of IP Quality of Service (QoS)
Services for JANET**

Lancaster / C&NLMAN Deliverable 1: Initial Findings

Editor: Christopher Edwards

Contributors: Mark Jameson, Michael Clarke, Ian Anderson, Christopher Allen

Version 0.3 ** FIRST ROUGH DRAFT **

4th July 2003

Contents

1	INTRODUCTION	2
2	INFRASTRUCTURE VIEW.....	3
3	THE NEED FOR QOS.....	4
3.1	COMPUTER SCIENCE AND DEPLOYMENT PERSPECTIVE	4
3.2	IDENTIFICATION OF “REAL” USERS.....	4
3.3	AN INITIAL APPLICATION TRIAL	5
4	PLANNED TIMETABLE.....	6
5	QOS MONITORING RESULTS.....	6
5.1	RESIDENT NETWORK (RESNET).....	7
5.2	PORT MIRRORING ON / OFF CAMPUS	7
6	DEVICE CONSIDERATIONS	8
6.1	QOS SUPPORT IN VLANS.....	10
7	RESULTS OF VENDOR DISCUSSIONS	10
7.1	EXTREME	10
7.2	CISCO	11
8	OPEN ISSUES.....	11
8.1	INFRASTRUCTURE ISSUES	11
8.2	GENERAL ISSUES	12
8.3	WIRELESS DOMAIN.....	12
9	CONCLUSION.....	13

1 Introduction

This is the first release of deliverable 1 from the Lancaster / C&NLMAN combination taking part in the JANET QoS Development Project. Getting the QoS work “off the ground” at Lancaster has been difficult due to the significant workloads that members of both the Information Systems Support (ISS) and Computing Departments at Lancaster University have been under. We have therefore agreed to release an initial version of the deliverable 1 now, and a second version of the document later in the summer. Whilst we have perhaps not made as much progress as would be ideal, we have taken a number of steps in terms of understanding the issues that QoS introduction will give us, the equipment that will be impacted, etc... along with an analysis of the QoS-marked traffic already existing on our networks.

As will be underlined within the “Infrastructure View” section of this deliverable, the networks governed by the Lancaster/C&NLMAN partnership are significant both in terms of size, reach, feature, and complexity. Upgrades to the infrastructure are continuing. Perhaps the most significant of these upgrades is the increase in the C&NLMAN backbone between Lancaster and SuperJANET from 155 Mbps to 10 Gig Ethernet. Other work includes the continuing addition of institutions to the CLEO network, and upgrading of a number of the main routers on the CLEO network. All of these upgrades of course add to what is already a complex scenario when considering the introduction of QoS mechanisms.

Given the complex infrastructure, we also have a number of options as to “where to start” in terms of QoS deployment. C&NLMAN and the University backbone seem the obvious place to start, but it has also become clear that perhaps the CLEO network is the more interesting network to explore given the range of networking technologies (both wired and wireless) in use, and it’s overriding “need” for QoS. From this perspective, we have decided to look at the introduction of QoS functionality from a number of perspectives, serving the C&NLMAN backbone and parts of the University campus network as an initial trial, and later parts of the CLEO network.

We have also had to be a little more realistic about the application domains that we are likely to address within the trial. Whilst we have significant interest at Lancaster (and the surrounding networks) in a number of different application domains (relevant in terms of QoS), support for VoIP and video conferencing is in high demand. Therefore, this is where we hope to start.

The remainder of this document is structured as follows; Section 2 has a brief look at the infrastructure that the Lancaster/C&NLMAN combination are in control of. Section 3 looks in a little more detail at the potential need for QoS support within the infrastructure. Section 4 looks at an initial timetable for the QoS work at Lancaster. Section 5 provides an overview of the results of initial QoS monitoring results carried out. Section 6 outlines the devices that require consideration within the infrastructure. Section 7 provides some notes relating to discussions that we have had with the main equipment vendors involved, Extreme and Cisco. Finally, section 8 describes a first shot at outlining some of the open issues that need to be addressed.

2 Infrastructure View

Figure 1 shows the topology of the CANLMAN backbone, the connection to the Lancaster University core router, and example connectivity down to end users in the Computing Department and Management School.




Figure 1: C&NLMAN and Lancaster University

Figure 2 shows the CLEO network. CLEO is made up of 155Mbps fixed and microwave links (plus wireless 10Mbps and 2Mbps) connecting Lancashire and Cumbria. As well as the different physical medium involved, there are also relay links where multiple sites share one 10Mbps connection.

JANET

CAN

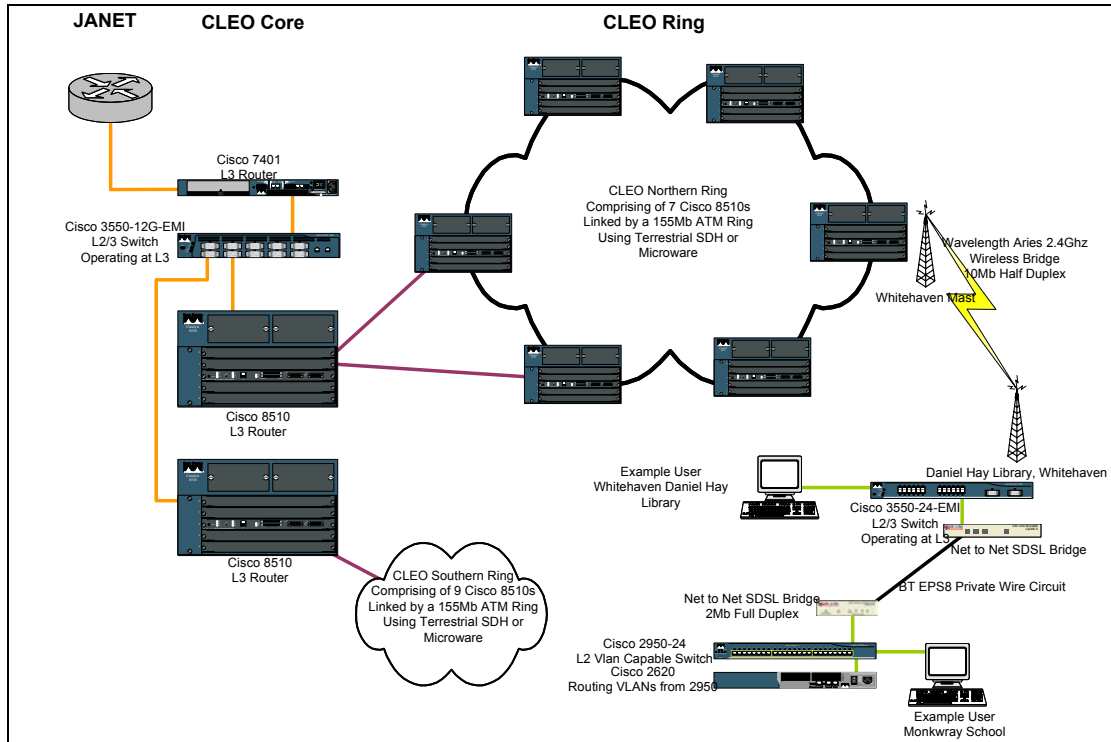


Figure 2: The CLEO Network

3 The Need for QoS

The need for QoS and the interest in the work that is taking place in the JANET QoS Development Project in general stems from a number of areas.

3.1 Computer Science and Deployment Perspective

The Computing Department at Lancaster University has for many years been involved with research into the area of QoS within distributed systems, and more recently its support within IP-based networks. Deploying QoS (“seeing it to the end”) is of great interest!!

As providers of a networking service, the ISS department obviously have an interest in practical networking, and a practical interest in providing support for QoS. There are also a number of services (e.g. the video conferencing suite) that would in the long term benefit from QoS provision.

3.2 Identification of “Real” Users

As a provider of network connectivity to many local and regional organisations, a number of “real” users of QoS technology have already been identified. For example:

- St Martins College have a MyTel VoIP system and ideally require QoS connectivity between the areas of Bowerham, Lancaster and Ambleside;
- Lake District National Park Authority are using VoIP. ISS provide a VPN from their main office down to the University. They are interested in QoS, and ISS

would obviously be interested in understanding how QoS support is provided within VPNs;

- CLEO have purchased a Video Conferencing H323 gatekeeper, for CLEO connected institutions (e.g. Schools) to book time on and use;
- Lancaster University are likely to require VoIP support at different points within the near future. The new LEC building is to have VoIP, and will the new InfoLab21 building. The plan initially is to have this supported by an entirely separate network (i.e. separate to the University backbone). Realistically, this work needs to be completed keeping in mind the fact that eventually we will be looking to support VoIP across the entire campus network. We also have the issue of supporting softphones against true IP phones; we understand that 1 softphone has the overhead of 20 true IP phones...

3.3 An Initial Application Trial

An initial application trial that we are hoping to complete will be to support a softphone application (ViaVideo) on a small number of machines, perhaps in the Computing Department. VidVideo is a “cheap and cheerful” application supporting IP network Quality of Service (QoS) with Polycom Video Error Concealment (PVEC) and DiffServ. As shown in section 2, the link between machines in the Computing Department and ISS goes through a Cisco 3548, an Extreme 7i, and from there it is VLANed / trunked up to the Extreme BlackDiamond. A screenshot from the ViaVideo application can be seen in Figure 3.



Figure 3: ViaVideo

4 Planned Timetable

Between now and the end of the calendar year, we intend to attempt the following:

- 1) Remark QoS traffic that we don't know about.
We are relatively sure that this could be done initially, although would need to be scheduled for "at risk" time;
- 2) Sorting the "known" marked traffic against the "unknown" marked traffic on the BAR could be processor intensive, although certainly not impossible. One option that we could try would be to move all the "known" qos clients into a small subnet range of 148.88.x.x then exclude those from the remarking access control list, for example:

```
<suedo config>

    interface gigheter 0/1
        description Uplink to canlman via this interface
        access-group outremarking out

    access-list extended outremarking
        permit ip 148.88.255.0 0.0.0.255 let_it_go_untouched
        permit ip host 148.88.0.0 0.0.255.255 remark_it

</suedo config>
```

- 3) QoS marking of packets running ViaVideo from a machine in Computing, where we would begin to honour the markings all the way back to ISS and eventually to C&NLMAN;

This can be broken down into a number of areas:

- i/ The client;
- ii/ the application;
- iii/ the "infrastructure"

What is described above seems like a reasonable initial target, without trying to take things too far, too quickly. We intend to define further plans as this work progresses, obviously subject to the success of the above.

5 QoS Monitoring Results

At this earlier stage in our consideration for QoS deployment, one of the initial targets we set was to try and understand how much marked traffic was traversing different components of the network. This was completed using a number of pieces of open source traffic monitor software.

5.1 Resident Network (RESNET)

One of the initial discoveries made was that 4%-8% of the traffic leaving the RESNET was marked as Premium!! At the time of writing, we are in the process of investigating this further.

5.2 Port Mirroring On / Off Campus

Using a Cisco 7401, we were able to mirror the incoming and outgoing ports on the on / off campus link. Using a self developed program, we were able to identify how much traffic was marked in some way.

500,000 IP packets were sniffed. Of these, 7029 (1.41%) contained data in the TOS field. A breakdown by TCP source and destination port can be found below in tables 1 and 2. Where the port numbers are well known, the application has been listed. As you will see, most of the traffic is on not-well-known ports. This will be in cases where applications have dynamically negotiated ports, which obviously makes identification somewhat harder. These are very crude measurements, but they are a start, and show quite clearly that marked traffic does exist. At the time of writing, we are attempting to use NBAR to provide further analysis.

TCP DPort	Packets	Total packets	Percentage
722	751	499845	0.15
www	678	499845	0.13
32824	550	499845	0.11
3481	536	499845	0.10
4662	533	499845	0.10
2752	339	499845	0.06
2887	164	499845	0.03
32838	162	499845	0.03
ssh	143	499845	0.02
1516	133	499845	0.02
3059	129	499845	0.02
1020	118	499845	0.02
1392	105	499845	0.02
1430	99	499845	0.01
2325	98	499845	0.01
2329	95	499845	0.01
1557	89	499845	0.01
1499	87	499845	0.01
1504	84	499845	0.01
3561	78	499845	0.01
35256	76	499845	0.01
61769	68	499845	0.01
1168	51	499845	0.01

Table 1: Destination Port

TCP SPort	Packets	Total packets	Percentage
722	789	499845	0.15
www	726	499845	0.14
36304	421	499845	0.08
telnet	402	499845	0.08
32824	386	499845	0.07
ssh	255	499845	0.05
32838	221	499845	0.04
2050	212	499845	0.04
3012	190	499845	0.03
3889	190	499845	0.03
1645	184	499845	0.03
2636	164	499845	0.03
35256	143	499845	0.02
3493	133	499845	0.02
3698	132	499845	0.02
3171	129	499845	0.02
4397	127	499845	0.02
8000	102	499845	0.02
2526	99	499845	0.01
1561	98	499845	0.01
3305	87	499845	0.01
1020	85	499845	0.01
2584	78	499845	0.01
3725	73	499845	0.01

Table 2: Source Port

6 Device Considerations

We have a significant amount of heterogeneous equipment spanning C&NLMAN, Lancaster University and CLEO. In each instance, we will have to consider what QoS support is available, and what the configurations for the devices might look like.

On Campus:

- Alcatel 6024s (soon to be replaced)
- 3com fmss (soon to be replaced)
- HP procurve 2650
- Cisco 2950s (of various types i.e. number and type of ports)
- Cisco 3550s (of various types i.e. number and type of ports)
- Extreme 48i
- Extreme 5i
- Extreme 7i
- Extreme BD

In CLEO:

- Cisco 2950s
- Cisco 2620s
- Cisco 3550 (emi's)
- Cisco 8510s
- Cisco 7401s

In C&NLMAN

- Extreme BD
- Cisco 3662
- Cisco 7500
- Cisco 7200

Tables 3 and 4 show some basic detail regarding the QoS support provided within the equipment that makes up the path from C&NLMAN to Computing (table 3) and the CLEO network (table 4).

Name	Description	OS Version	Interfaces	QoS Support
Extreme Black Diamond	The Extreme BD is a Layer 3 switch that provides Lancaster University core routing.	6.2.1	Gigabit Ethernet, Fast Ethernet	DSCP, Layer 1-4 Classification, Priority Queuing, RED, Policing
Extreme Summit 7i	The Extreme 7i provides routing for the Engineering Department.	6.2.1	Gigabit Ethernet	DSCP, Layer 1-4 Classification, Priority Queuing, RED, Policing
Extreme Summit 5i	The Extreme 5i act as a Layer 2 switch for Engineering.	6.1.4	Gigabit Ethernet	DSCP, Layer 1-4 Classification, Priority Queuing, RED, Policing
Cisco 3548 XL	Provide connectivity to end-users.	12.0(5.3)WC(1)	Fast Ethernet	CoS Classification, Low/High Priority queuing

Table 3: QoS Support within Path Down to Computing Department

Name	Description	OS Version	Interfaces	QoS Support
Cisco 7401	The Cisco 7401 is the CLEO BAR. Provides access to C&NLMAN.	12.2(2)DD3	Gigabit Ethernet	DSCP, CBWFQ, Policing, LLQ
Catalyst 8510 MSR	Provide backbone routing for the CLEO network.	11.2(1.4.WA3.0)	Gigabit Ethernet, Fast Ethernet, ATM	?
Catalyst 2950	Example of equipment providing connectivity to end-users	12.0(5.3)WC(1)	Fast Ethernet	(4 egress queues) Strict Priority (L2 CoS) WRR

Table 4: QoS Support within CLEO

6.1 QoS Support in VLANs

Within our infrastructure we use a significant amount of VLAN tagging (802.1Q). One of our concerns in terms of QoS was whether or not the equipment would support / maintain / honour any layer 3 QoS provision.

The following text was found in:

<http://www.tml.hut.fi/Opinnot/Tik-110.551/1999/papers/08IEEE802.1QosInMAC/qos.html>

“No Quality of Service mechanisms are defined in this standard, but an important requirement for providing QoS is included in this standard, e.g. ability to regenerate user priority of received frames using priority information contained in the frame and the User Priority Regeneration Table for the reception Port.”

Also, Cisco documentation at:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s17/12s_tos.htm#wp1021949

“configure the ToS and Time-to-Live (TTL) byte values in the encapsulating IP header of tunnel packets for an IP tunnel interface on a router. The Tunnel ToS feature is supported on Cisco Express Forwarding (CEF), fast switching, and process switching forwarding modes.”

All of the above seems to confirm that we should be able to honour our QoS markings with the VLANs.

7 Results of Vendor Discussions

As part of the first steps towards deployment of QoS technologies at Lancaster, we met with the two main equipment vendors that together provide a large proportion of the equipment used across the C&NLMAN, Lancaster University and CLEO networks.

From a high level perspective, the architectural difference between the two vendors is significant. Extreme appear to use a common architecture across the range of equipment that we use, meaning that in theory, QoS provision should behave in a similar way whichever model is in use. With Cisco, QoS support appears to depend very much on particulars of the individual device; model, memory capability, divide between hardware and software, etc...

7.1 Extreme

The main focus at Extreme seems always to have been on the support of packet marking both at layer 2 (802.11p) and layer 3 (DiffServ). There is a clear mapping between the two. In terms of performance, all of the packet processing is carried out in hardware, even if we go as far as looking at a packets “port” specification. This in theory should mean that performance hits for using QoS functionality should hopefully be minimal.

Everything appears to be governed by the policy / access list defined. Hence, DiffServ is automatically switched on, but without a defined policy / profile, nothing will happen to marked traffic. The general rule seems to be to create a policy

framework that understands what is coming in at layer 2, and decides what to do with it.

Sample configurations have been provided by Extreme, but have not been included in this version of the document.

7.2 Cisco

Cisco classification can work in the following ways:

- all traffic on a port;
- IP precedence;
- MAC address;
- IP access list (standard and extended);
- NBAR – a potentially CPU intensive application that allows admins to write access lists based on application.

Queuing mechanisms supported are as follows:

- Priority queuing;
- Custom queuing;
- WFQ (the default QoS mechanism, probably already running);
- Class-based WFQ, with priority queue;
- RED;
- Shaped Round Robin coming soon (might be useful for LBE tests).

GSRs etc have fixed hardware queues. This means a lack of flexibility, but very high speed. Cisco platforms supporting CEF mean that the addition of DSCPs and the forwarding based on the DSCP does not impact on forwarding speed. 2950, 3550, 6500 all set and honour DSCP values in this way.

Things that we have been told to look out for:

- Make sure that we understand what the specific platform is that we are using;
- Realize that there may be some bugs;
- Understand the traffic on the network;
- Be aware that some queuing algorithms may not given exactly what is expected;
- Start small.

Sample configurations have been provided by Cisco, but have not been included in this version of the document.

8 Open Issues

We are in the early stages of deploying QoS technologies within our infrastructure. We thought it would be useful to round off this document with an overview of some of the issues we have come across through research, discussion and experimentation. We hope that (where required) we will slowly be able to provide solutions to each;

8.1 Infrastructure Issues

We have come across a number of “infrastructure issues” during our work.

- At a general level, the migration to Ethernet (and it's high speed derivatives) into the WAN to replace the more traditional telco derived technologies (such as ATM, SDH, SONET which have "qos" type mechanisms built in) extends the QoS problem to a wider scale
- Additional equipment is to be introduced at Lancaster within the next 4 months. A Cisco 6506 is being introduced to help take some of the load off the BlackDiamond, particularly at layer 3. Access control functionality is failing to scale to our requirements
- Large scale network upgrades can include:
 - o Changing the configurations (on every device?);
 - o Possibly upgrading firmware (often problematic in itself - break other things/require upgrades etc);
 - o Older hardware has to be replaced;
 - o Clever (and potentially processor intensive/dangerous) acls/inspection required at many places - including our stretched BARs;
 - o Client machines and applications have to do the right thing.
- Policing, and policies to decide how to police are required. The way these policies are set and monitored needs careful consideration. Who / what detects whether the policies are doing the right job? A full time post?! This also has the potential to touch on the issue of intrusion into users private affairs.
- Inter-vender issues - do all vendors interpret and honour the bits the same way? Different configurations will be required on different platforms (by same vendor). Multiple vendors might make things even harder. Honouring the bits at every hop (obviously) includes:-
 - o VLAN trunk hops (where we're effectively multiplexing over a shared link)
 - o Tunnels (where packets get put in a "wrapper" - can we inspect and honour the bits inside embedded in that wrapper)

8.2 General Issues

Some more general issues that we have discovered are as follows:

- Problems with unstable code of open source traffic monitoring software;
- Small scale upgrades required for testing, e.g. for Cisco NBAR.

8.3 Wireless Domain

Wireless domain presents a number of interesting challenges for us:

- 1/ By it's very nature it is less reliable than a wire - leading to:
 - a/ we loose more data (in some strange and interesting ways);
 - b/ there's more variation in delivery time for data;

c/ "inter-layer problems" cause errors to compound (including duplicates). Bad news for some TCP sessions.

2/ the unlicensed 2.4Ghz radios we use are 11Mb 1/2 duplex, so relatively small bandwidths;

3/ we sometimes deploy the radios in point to multipoint mode, so we effectively have a shared medium but with all the added complexity of radio communications;

4/ the radios do not currently support any type of qos or tos - though we have good relations with the equipment vendors, they may be able to help here;

5/ currently the radios pull small packets off the wire and keep them buffered until they have enough to justify sending out across the airwaves (ostensibly for better bandwidth utilization). This causes interesting delay characteristics;

6/ Due to local geography we often have to string chains of these radios together to reach areas of Cumbria;

7/ because of point 6 we often have no choice but to hang a number of "customers" off a single link - leading to contention for resources.

Despite all of this, CLEO is a very successful operational network providing connectivity to rural regions that without wireless capabilities would find it difficult and costly to gain access to the Internet at any sort of reasonable speed.

9 Conclusion

We have a long way to go, but have a much better understanding as to what it is we are up against!