JANET QoS Development Project

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

Lancaster / C&NLMAN Deliverable 2: Test Results

Editor: Christopher Edwards

Contributors: Ian Anderson, Mark Jameson

Version 1.0

8th October 2004

Contents

1 Introduction

This document is the second deliverable from the Lancaster / C&NLMAN combination taking part in the JANET QoS Development Project.

The aim of this work is to attempt to validate the possibility of deploying QoS mechanisms within the Lancaster / C&NLMAN infrastructure, in line with the JANET QoS Development "Technical Framework" document [1].

The first deliverable [2], released in July 2003 outlined our initial findings in the area. It outlined our perspective on the need for QoS support within our networks, and a view of the infrastructure. It also described some considerations with regard to the devices used within the network, and some discussions that had taken place with the main equipment vendors, Extreme and Cisco. The document finished by describing a number of "open issues".

Having set the scene in the first deliverable, this second document will report our findings. Section 2 provides an updated view of the Lancaster / C&NLMAN network infrastructure, including details as to how the network has been instrumented to allow the QoS trials to take place. Section 3 provides an updated version of the QoS Monitoring Results that took place during the early stages of the project, that outline the traffic found on the University campus network that had packets marked with QoS values. In the following 4 sections (sections 4 – 7) we discuss the different trials that took place as part of the trials. In each instance, we outline the aims of the trial, the results, and discuss our findings. In section 8 we provide a number of conclusions regarding the work completed.

2 Infrastructure View

Figure 1 shows the basic topology that we are dealing with at Lancaster, showing the CANLMAN backbone, its relationship to the JANET infrastructure, and the connection to the Lancaster University core router. It also shows example connectivity down to end users in the Computing Department and Management School, and provides details of the link speeds available.

In taking part in the JANET QoS trials, we were keen to make as much use of the existing infrastructure in order to ensure that the trials were as “real” as possible. As a result, we did not attempt to add any additional devices into the data path that the QoS traffic would encounter. We also stayed away from the idea of replicating any of the infrastructure in order to run the tests.

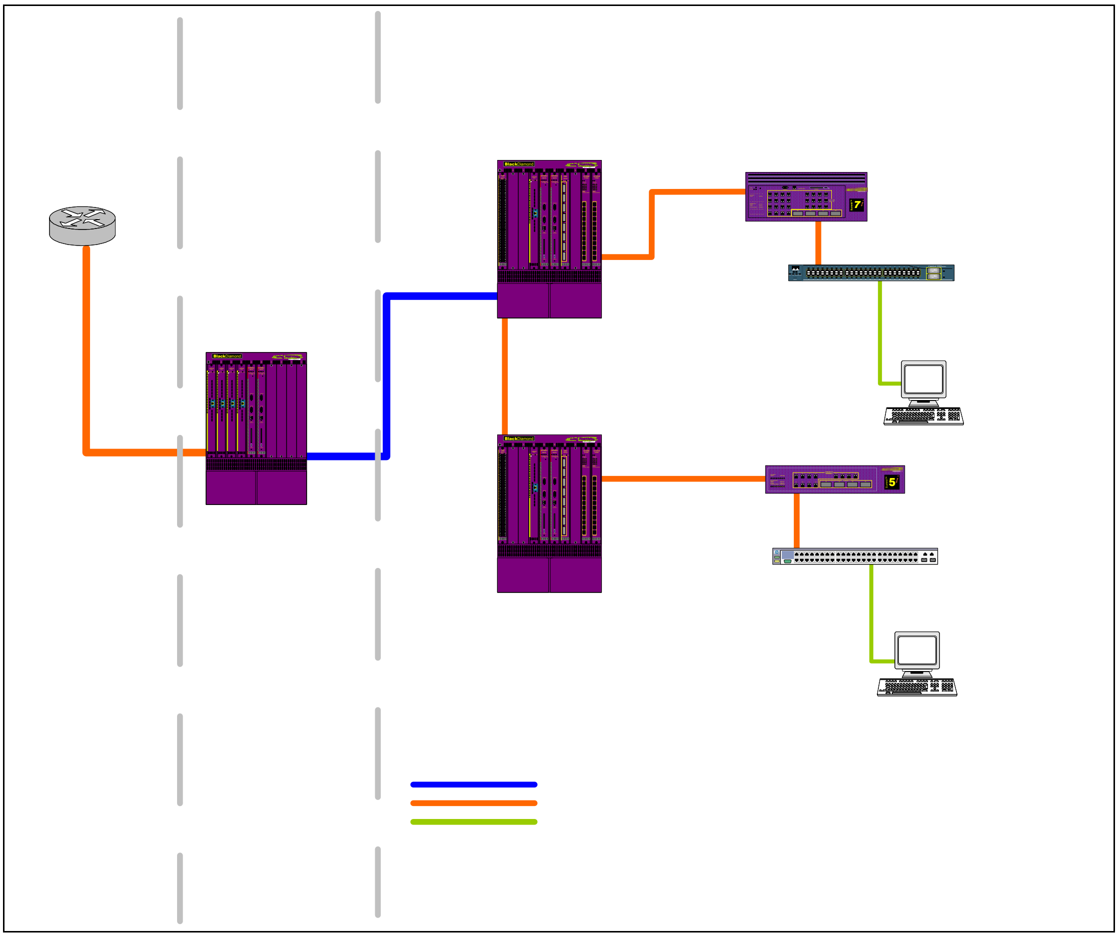

Figure 1: C&NLMAN and Lancaster University

2.1 Instrumenting the Infrastructure

Figure 2 provides a more detailed illustration of the network at the time we started the tests, and shows the means by which we instrumented the network in order to carry out the tests.

The main features of this instrumentation are as follows:

- Enabled QoS on the devices on the path down to the SECAMS/Computing building
- Added 2 iperf servers (iperf-edge and iperf-core). These devices were used to load the external link from the University. We located them at different links on the campus network in order to allow for localised sourcing and syncing of traffic for any “on campus only” experiments
- Two Cisco 7940 phones were added to the infrastructure, again at different points in the network
- A Cisco 805 router was hung off the Lancaster University core Black Diamond switch to act as an SAA agent

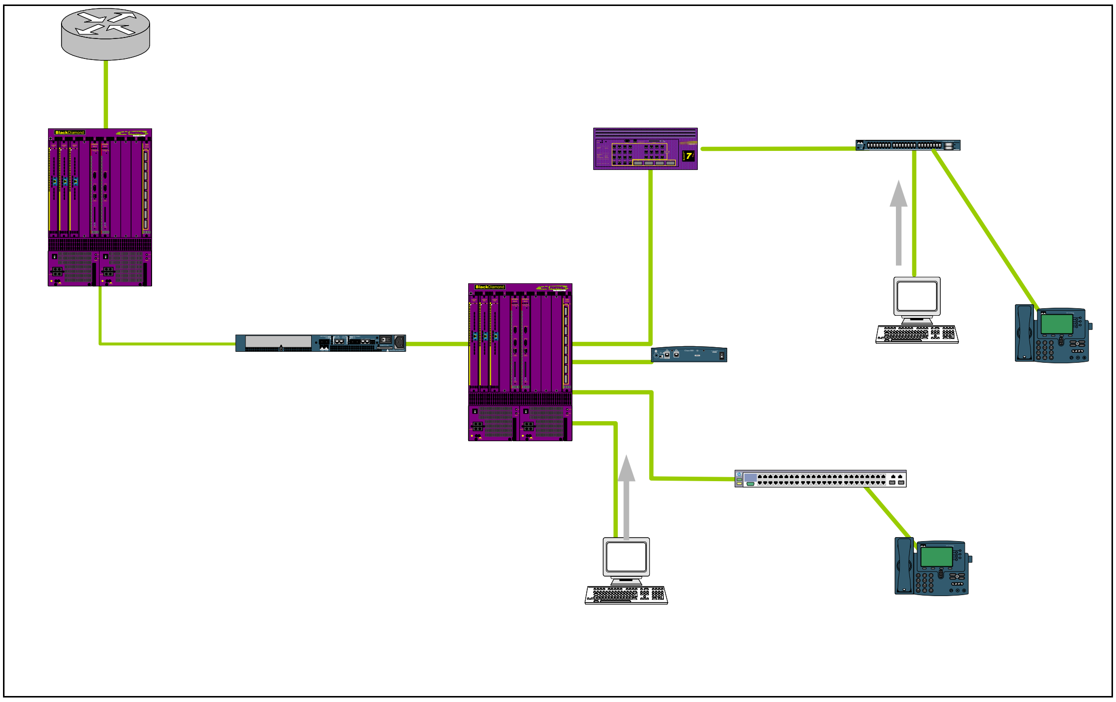

Figure 2: Network Status at Point of Initial Testing

Throughout the duration of the trials, a number of changes to the infrastructure were made:

- The Cisco 805 acting as the SAA agent was replaced by a Cisco 1721 due to the fact that we were seeing a number of strange graphing results. This seemed to be down to the performance of the Cisco 805, and is discussed further in section X. Other partners followed suite.
- The Lancaster University BAR (a Cisco 7401 in figure 2) was replaced by a Juniper M7i between the tests in March and July. The new Juniper device was configured to provide the same QoS provision as the original Cisco

device. Appendix 2 shows the configuration for the original Cisco 7401. Appendix 3 shows the configuration for the new Juniper.

3 QoS Monitoring Results

Prior to the main testing phases that required booked “network at risk” sessions (discussed in sections 4 – 7), one of the initial targets 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 monitoring software, and an application developed by a member of the network support team at Lancaster. The source code for this can be found in Appendix 4.

Using a Cisco 7401, we were able to mirror the incoming and outgoing ports on the on / off campus link. Using the program developed internally (see Appendix 4), we were able to identify how much traffic was marked.

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.

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

4 March 2nd Tests

4.1 Aims

In this the first inter-partner test, we were working with Manchester to see if we could impact on the quality of the traffic between sites. The main parameters that we were concerned with were RTT and packet loss. The overall scope for the experiment was to see how much traffic was required on the various links in order for us to see a degradation of quality both from the measurement points in place, and any voice calls placed between the two sites.

The combination of Manchester and Lancaster was interesting; Manchester were using a test network, whereas Lancaster were loading their main JANET access link. Further details of the Manchester configuration can be found in [3].

Note that:

- Whereas Manchester were saturating their network link by bringing three traffic classes to their defined limits (20% for Premium, 75% for BE and 5% LBE), Lancaster were loading purely with BE traffic

- For this first trial, network saturation was achieved using (BE) TCP traffic

- The measured maximum throughput of the link was 140 Mbps

- Configurations for the C&NLMAN and LU Bar devices can be found in Appendix 1 and Appendix 2 respectively (although as noted above, the traffic loading the network was BE)

- Objective measurements were obtained via the RTT data from the SAA. From the Lancaster perspective, the IP phone statistics did not really help; Manchester had more success in this area

- Subjective measurements were taken by monitoring the call quality on the VoIP phones

- Within this first phase of testing, all VoIP calls were made on a Best Effort basis

4.2 Results

This section outlines the results from the first set of tests. It begins with a timeline of the main events that took place, and is followed by a description of the results seen.

07:45	Initial call with Manchester prior to congestion. Both phones operating as BE. The call quality was perfect.
07:56	We added traffic to the scenario from two sources, iperf-core and iperf-edge using iperf (TCP, Best Effort). The traffic was destined for Manchester's traffic sync. The overall load out of JANET was 140 Mbps.
08:05	A further BE call was placed from Lancaster to Manchester. We observed a maximum jitter of 6ms (varying between 4ms and 6ms), with no packet loss. The call quality seemed unaffected.
08:10	Whilst Manchester changed the direction of traffic flow across their test infrastructure, we kept things constant.
08:11	We then placed a further BE call to Manchester. Whilst Manchester recorded a degraded call quality (and saw packet loss figures), we had no packet loss, and maintained a perfect call quality. It was our belief that this was the case because we were all congesting the links in the same direction, and therefore it was possible that the traffic received by Lancaster was not competing with any loading.
08:20	Manchester reversed the direction of their loading internally
08:38	Final call with Manchester, still with very high traffic loads, call quality seemed very good, although Manchester recorded packet loss figures.
08:45	First trial ended.

Table 3: Timeline for First Tests

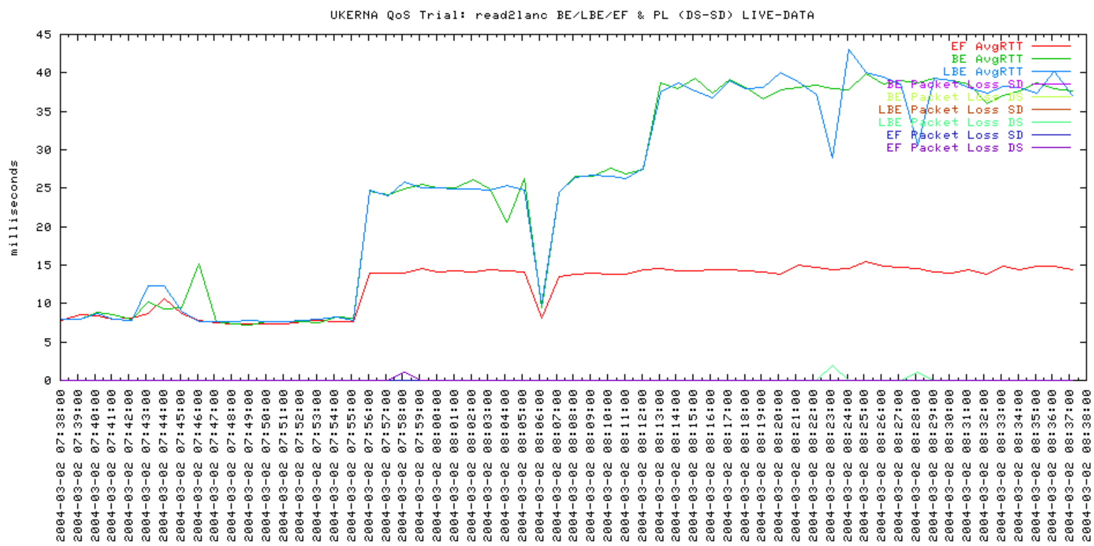

Figure 3: Reading to Lancaster

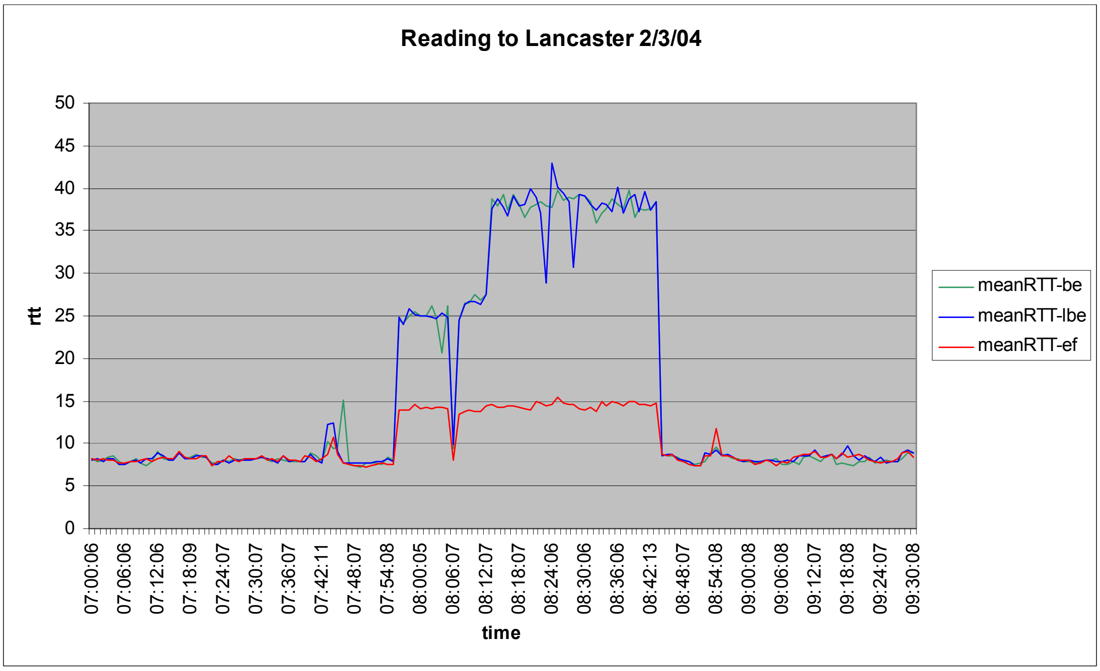

Figure 4: Reading to Lancaster

Figures 3 and 4 are showing the RTT times for the probes being sent and collected by the SAA devices. Figure 3 is generated by the monitoring infrastructure automatically, Figure 4 shoes a raw data re-plot. Both figures show clearly that as the traffic loading is ramped up at the beginning of the test, the RTT times for both BE and LBE increased significantly, whilst the EF RTT only increased slightly. Whilst further on into the trial the RTT for BE and LBE increased even further (up to 40ms), the EF RTT remained constant at around 15ms.

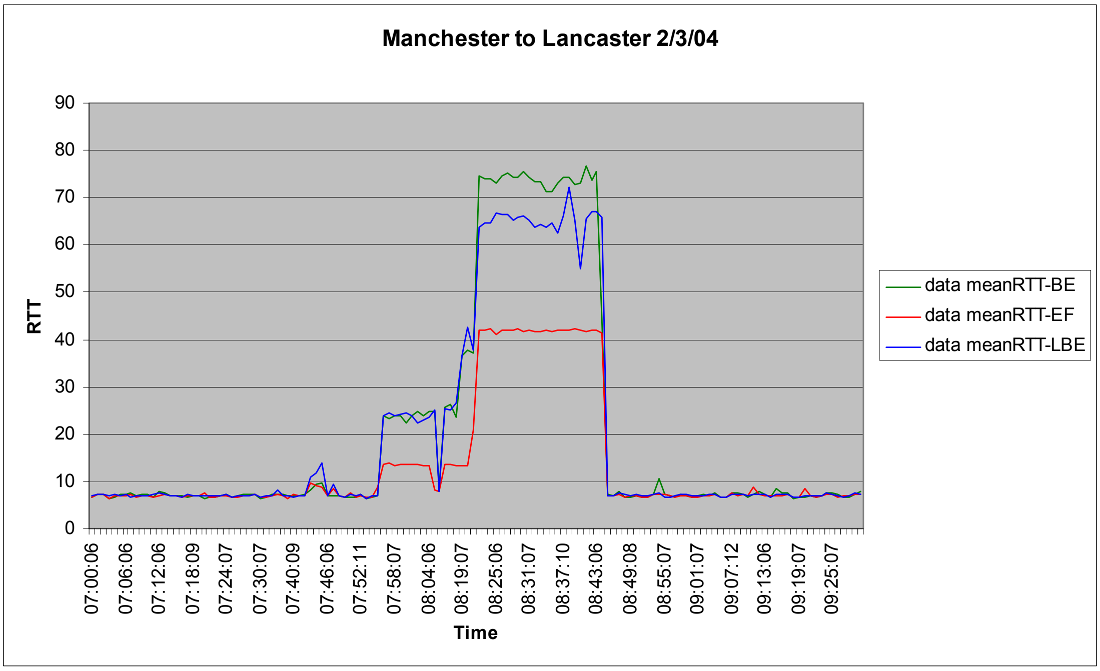

Figure 5: Manchester to Lancaster

Figure 5 shows a plot of results of the RTT times between Manchester and Lancaster, i.e. from the two MANs involved in the tests. When the traffic loading started before 08:00, the results looked similar to the earlier figures, in that we saw an increased RTT for BE and LBE, but less so for EF. However, at around the 08:20 mark, you can see that the RTT times for all classes significantly increase, including the figure for EF. We believe that this anomaly ties in with the packet loss figures discussed in the Manchester results [3], and the reversing of their internal traffic.

4.3 Conclusions

The main conclusions from the first tests were as follows:

- We were surprised by the speed at which the time past for the first trial, but still felt we achieved some reasonably valid results
- From the call perspective everything from the Lancaster end seemed fine, even when the network was heavily congested. From the Manchester side, even with some clipping on the audio, every word was comprehendible
- Other than the anomaly in figure 5, everything seemed to be working as expected, in that the EF RTT times were unaffected by the loading
- The fact that we were using a combination of BE and TCP for the traffic loading may have had a significant impact on our results. It is entirely possible that there were premium RTR UDP probes in the network, TCP's own congestion control was backing off to allow the premium through

5 March 9th Tests

5.1 Aims

The main aim of the second inter-partner test was to continue on from where we left off. During the first test we'd not got as far as placing a premium phone call, and we'd used only TCP traffic to load the links. We also wanted to see the impact of using UDP for link loading. The Lancaster infrastructure remained the same for the second test. Manchester made use of a gigabit test link after their issues within premium traffic previously.

5.2 Results

As with the last set of results, this section shows a timeline of the main events that took place, and is followed by a description of the results seen.

07:30	Initial base call with Manchester was made, with no QoS for the call, and no congestion. The call quality was perfect.
07:36	We started to saturate the 155Mbps link, with BE TCP traffic out of our two iperf boxes, iperf-core and iperf-edge. A further best effort call was made between Lancaster and Manchester, and at this point Manchester identified whole words and phrases being lost. This implies that the conditions on the Lancaster link were having a significant impact, although from our perspective, the quality of the call was again good.
07:50	We changed our loading traffic from TCP to UDP. We started the UDP traffic at a fairly low level, until settling on 50 Mbps from each iperf box, which combined with the background traffic meant that the link was fairly saturated.
08:08	At this point with links at both partner sites heavily loaded, we started to see some strange signalling behaviour, and actually found it difficult to establish a call at all (this will be discussed in further detail in section 5.2.1). Once a call was established, call quality was good for Lancaster, but there were crackles evident at the Manchester end, and their phone reported packet loss.
08:14	We then further ramped up the UDP test traffic to well in excess of the 155Mbps that our access link can support. We then had the situation where the JANET BAR was having to drop as we effectively have a gigabit pipe feeding into a 155 external link.
08:27	The BE call between Manchester and Lancaster was breaking up, and both phones were now recording packet loss. Both partners then reclassified the traffic from one VoIP phone to premium, and made a call. The call quality was perfect, with no packet loss recorded. The BE calls were unable to establish a call.
08:34	Lancaster lost external connectivity (this is discussed further in section 5.2.2). The test ended.

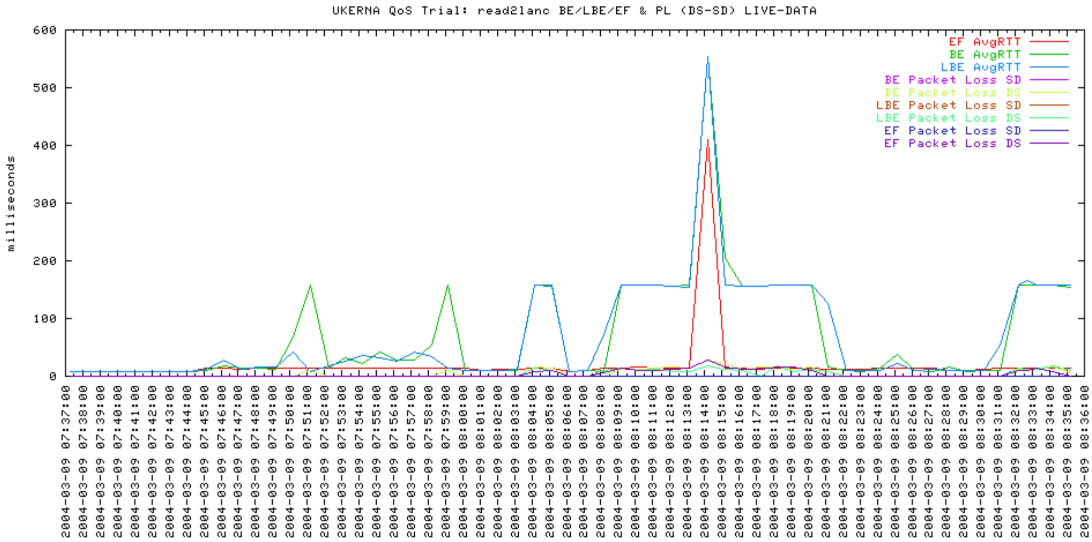

Figure 6: Reading to Lancaster

The graphs for this second set of inter-partner tests are certainly not as clear as the first set. Figure 6 shows the RTT times for the different classes, on the Reading to Lancaster path. Note that the blip at around 08:13 (when we introduced significant load into the network) has caused the scaling of the graph to be significantly different to that shown in Figure 3. What is clear is that (other than the 08:13 blip) the EF RTT remains low.

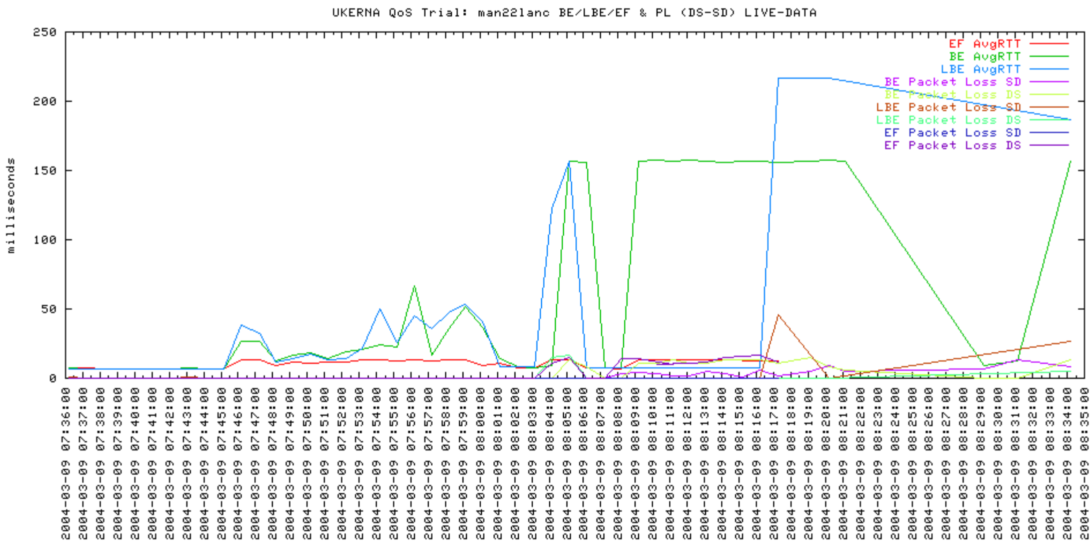

Figure 7: Manchester to Lancaster

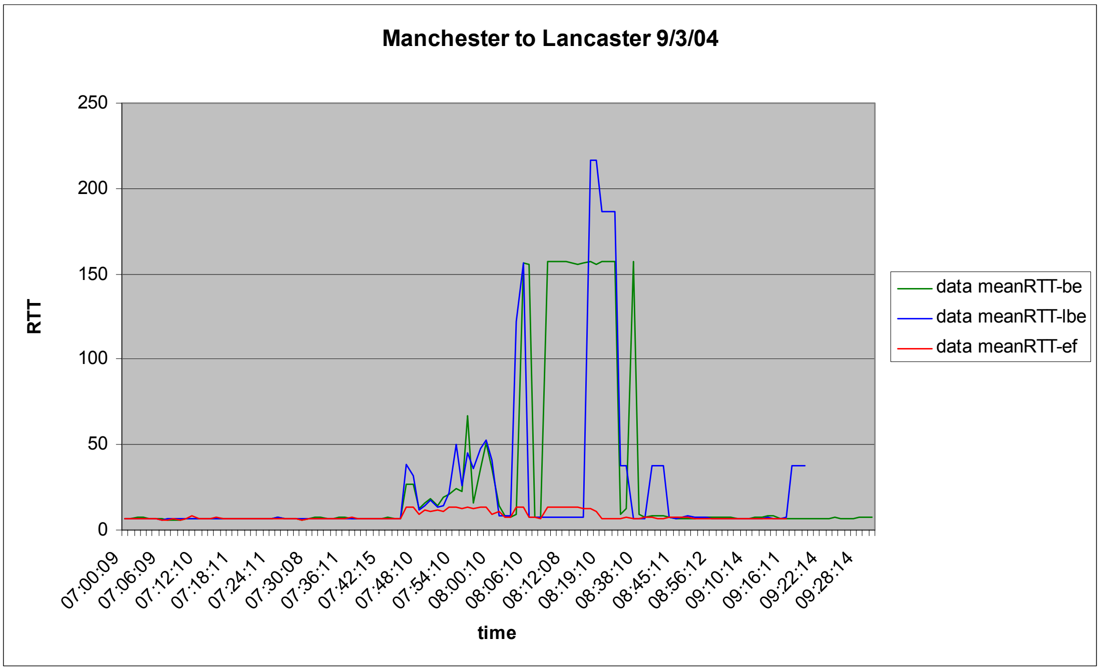

Figure 8: Manchester to Lancaster

Figures 7 and 8 show the same RTT times for the Manchester to Lancaster link, but Figure 8 is over a longer time period. Here we see a rather strange set of results, with some very high RTT figures for the BE and LBE classes. What again is clear is that the EF RTT times remain low.

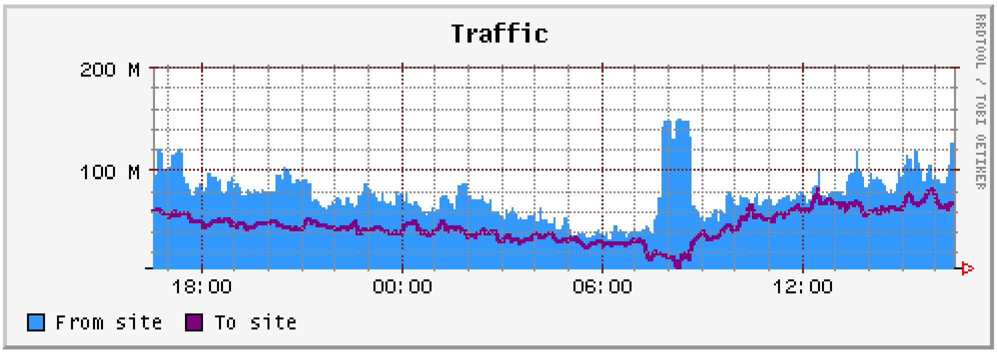

Figure 9: Link Traffic for Warrington to Lancaster BAR

Figure 9 is plotted from the JANET Netsight figures, and shows that clearly the Lancaster 155Mbps link reached saturation point during the test. This raised an interesting point because the Lancaster BAR is a JANET device, we were unable to query the device to see the impact that the loading was having. This gave us a very “real world” view as to where inter-domain rules are likely to be applied and policed.

5.2.1 Signalling Loss

As BE load was applied to the link(s), the call quality to Manchester became progressively worse with entire sentences dropping from the conversation and jitter values of 400+ being recorded on the phone. The graphs from the SAA monitoring appeared to be doing exactly the right thing with the RTT for EF traffic remaining low with the RTT for BE and LBE increasing. At this point with the phones being on BE and the link heavily loaded, as discussed in the timeline we began to experience some problems with the phones themselves. If we picked them up there would be no dialtone, and on 2 occasions at Lancaster we answered a call from Manchester only to have the phone continue to ring! There are some potentially interesting call manager problems here. Once the call is in progress the traffic flows directly between the two handsets, but the signalling that handles the set-up and clear-down of the call is handled from the call manager. Therefore if any load or adverse conditions are placed on the link affecting the call manager then call setup could be affected and other factors such as delayed dialtone could be experienced.

5.2.2 Availability of an Additional Link?

As we began to really ramp up the traffic during this second test, we realized that there was an additional link in the equation that we had neglected to factor in. Figure 10 shows the second link via the Lancaster BAR2.

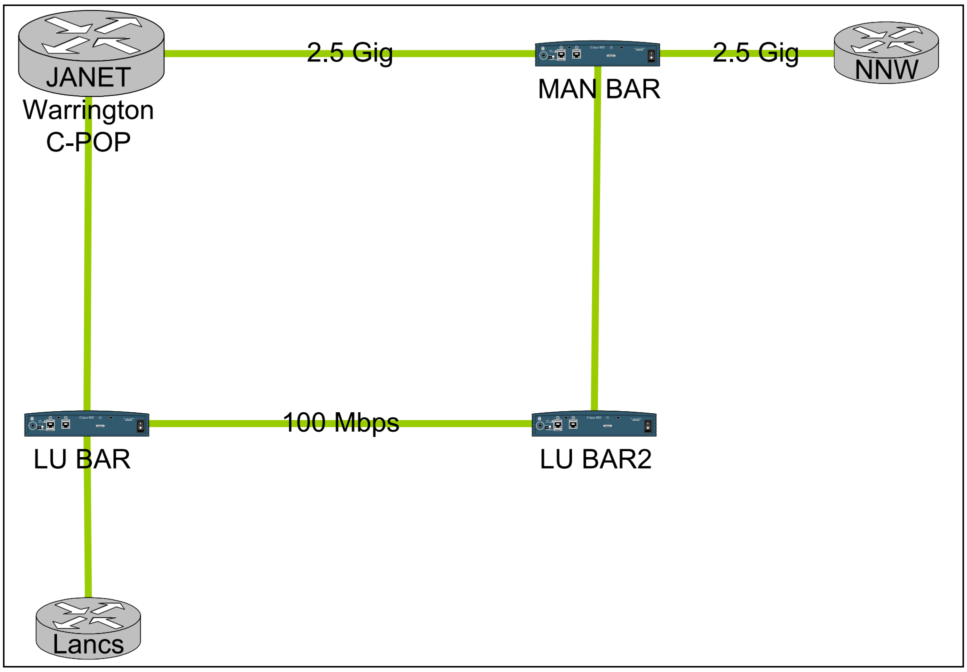

Figure 10: Additional Link Causing Complications

The additional 155Mbps link is for the RBC interconnect to which CLEO are connected. As the CANLMAN 155Mbps link approached saturation, it is possible that the traffic fell over (possibly at the BGP level) to also use the 155Mbps link via Manchester. This will no doubt have caused a few interesting asymmetric routing conditions for the traffic flows.

We believed that we had managed to completely congest both of the links, with the levels of traffic leaving CANLMAN approaching 230Mbps, although we were unable to monitor the loading on these links. The lack of visibility of the loading of the links was a problem, as we were unable to be 100% sure about our findings. A call from JANET Operations at about 08:30 to saying “we’ve lost C&NLMAN” is a clear indication that we were doing a good job of loading the link!

5.3 Conclusions

The main conclusions from the second tests were as follows:

- From the perspective of the VoIP calls, Manchester continued to have poor call experiences, although Lancaster reported significant packet loss as the traffic loads increased

- It was clear that we were loading up the BAR beyond capacity, and that there is a possibility that we were “forcing” traffic down an additional link

- Using UDP for our traffic generation seemed to be a lot more effective in terms of loading the link, but meant that we got limited feeback; this was a particular issue from the perspective of trying to understand what was happening at the Lancaster JANET BAR

- Some high RTT “blips” on the graphs made them a little difficult to follow

- EF calls worked!!

6 March 16th Tests

6.1 Aims

The third test was a little more low-key than the previous two, where we were really trying to replicate some of the behavior seen in the second test, whilst attempting to discover if the additional link was playing a role in carrying the traffic once the network became overloaded. It was deemed as important that we found an answer to this before progressing any further, in order to avoid confusion.

Our rather “basic” aims meant that we could carry out these tests without the interactive help of Manchester. They used the opportunity to carry out some testing with UKERNA, whilst allowing us to continue to send our iperf traffic towards them.

We loaded the networking using UDP BE traffic from the two iperf boxes iperf-core and iperf-edge. We gradually ramped up the traffic at set time intervals, beyond 155Mbps.

6.2 Results

Whilst increasing the amount of background traffic on the network, we carried out a series of traceroutes to determine if the traffic went down a different path, and it didn’t. This is despite carrying out similar traceroutes the week before, where the results seemed to indicate that both links were being used, giving an aggregate bandwidth of 155Mbps+100Mbps.

In conjunction with the tests, we got in touch with the JANET NOSC in order to determine their thoughts on the situation:

- They should be configured such that all traffic to and from C&NLMAN uses the 155Mbps to Warrington and all traffic to and from CLEO (who connect to lancaster-bar2) uses the 155Mbps to Manchester
- C&NLMAN traffic should only use the 155Mbps to Manchester if the line to Warrington goes down or the routing across that link fails (and similarly for the CLEO traffic)
- The JANET NOSC could see no sign of either having happened during the at-risk period. During normal circumstances, the only traffic traversing the 100Mbps link should be traffic between CLEO and C&NLMAN
- There should certainly be no load-sharing

6.3 Conclusions

The main conclusions from the third tests were as follows:

- It appears that the theory of using the second link may have been incorrect, although this does not explain some of the peculiarities seen during the previous weeks tests
- If the second link is not being used, this confirms that the bottleneck is at the JANET BAR
- The QoS probes were unable to return any useful values during this test

7 July Tests

7.1 Aims

Despite the main trial period ending in March, at Lancaster we decided that we would carry out an additional set of tests in order to try to gauge the impact of a number of changes to the network infrastructure. The main changes were:

- The Cisco 7401 (labelled LU BAR in Figure 2) was replaced by a Juniper M7i
- The JANET BAR was changed from a Cisco 7500 to a Cisco 12406
- The Cisco 805 used as the SAA Agent was replaced by a Cisco 1721

Note that our JANET link speed remained at 155 Mbps. The main aim was to ensure that the new equipment performed in the way we expected. Of particular interest was

the Juniper device, as this was the first of its kind to be introduced into our infrastructure. The Juniper configuration can be found in Appendix 3.
We loaded the networking using UDP BE traffic from the two iperf boxes iperf-core and iperf-edge. We gradually ramped up the traffic at set time intervals, up to 200 Mbps.

7.2 Results

Figure 11 shows that we had significant traffic running across the new Lancaster BAR.

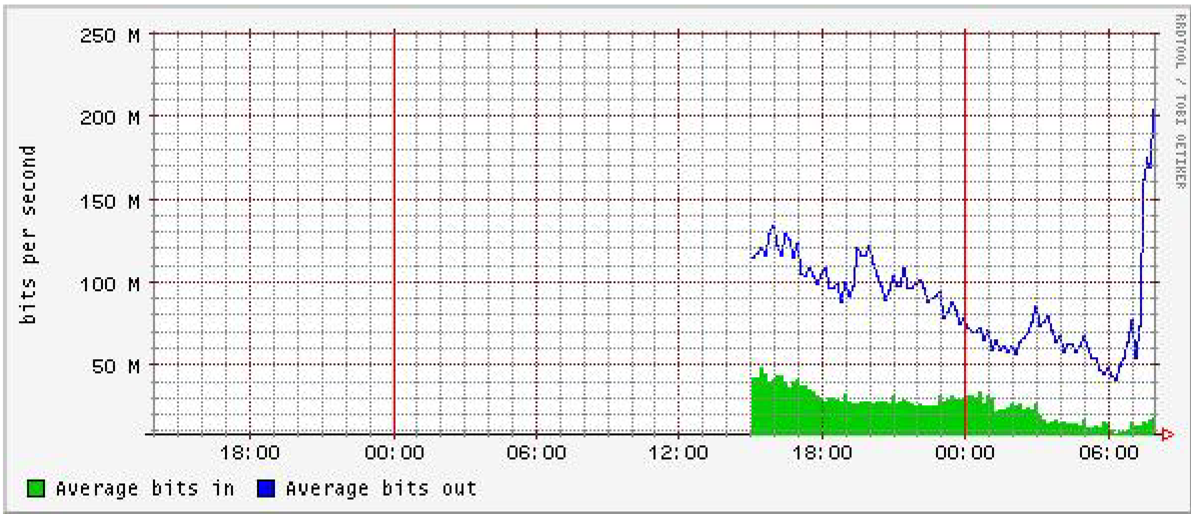

Figure 11: Traffic Across the Lancaster BAR

Figure 12 shows what was expected in terms of the RTT times during the period when the network was congested. Between 07:32 and 07:46 the RTT times for LBE and BE increased from 7 ms to 14 ms, where as the time for EF probes increased from 7 ms to 10 ms.

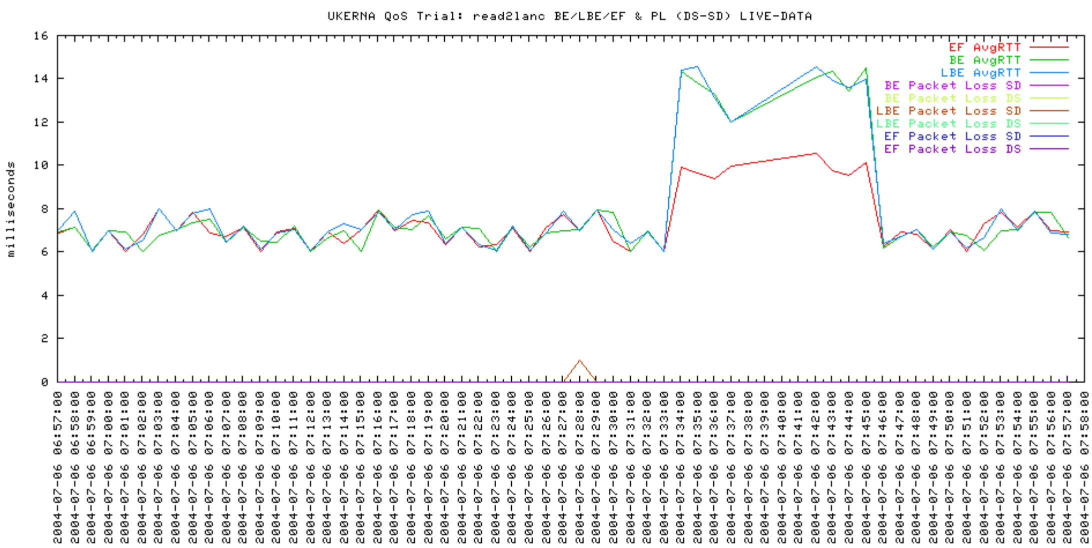

Figure 12: Reading to Lancaster

Despite seeing fairly expected results from Figures 11 and 12, during the tests we carried out a number of traceroutes in order to ensure that everything was as expected. During the point of loading some unexpected routing took place between Lancaster and Net North West. Compare the traceroutes in Figure 13. The address 146.97.35.213 resolves to a gig interface on the Warrington C-POP. The paths taken are shown in Figure 14. It is not clear as to whether these inconsistencies were due to other activities taking place during the “at risk” time, although we were unaware of any at the time.

```
[andersoi@tea andersoi]$ /usr/sbin/traceroute 194.66.26.42
traceroute to 194.66.26.42 (194.66.26.42), 30 hops max, 38 byte packets
 1  jips (148.88.8.6)  0.910 ms  0.569 ms  0.534 ms
 2  gw-lancs (194.80.32.1)  0.685 ms  0.783 ms  0.565 ms
 3  194.81.46.1 (194.81.46.1)  0.896 ms  0.936 ms  0.806 ms
 4  lanc-bar1.ja.net (146.97.40.9)  0.882 ms  0.789 ms  0.774 ms
 5  146.97.35.213 (146.97.35.213)  3.061 ms  3.128 ms  2.950 ms
 6  po1-0.manchester-bar.ja.net (146.97.35.166)  3.540 ms  3.577 ms  3.493 m
 7  gw-nnw.core.netnw.net.uk (146.97.40.202)  3.670 ms  3.828 ms  3.776 ms
 8  beacon1.netnw.net.uk (194.66.26.42)  3.714 ms  3.698 ms  3.663 ms



[andersoi@tea andersoi]$ /usr/sbin/traceroute 194.66.26.42
traceroute to 194.66.26.42 (194.66.26.42), 30 hops max, 38 byte packets
 1  jips (148.88.8.6)  0.705 ms  0.566 ms  0.616 ms
 2  gw-lancs (194.80.32.1)  0.671 ms  0.588 ms  0.571 ms
 3  194.81.46.1 (194.81.46.1)  1.134 ms  0.966 ms  0.789 ms
 4  lanc-bar1.ja.net (146.97.40.9)  0.939 ms  1.014 ms  1.020 ms
 5  manchester-bar.ja.net (146.97.37.57)  2.358 ms  2.337 ms  2.223 ms
 6  gw-nnw.core.netnw.net.uk (146.97.40.202)  2.615 ms  2.461 ms  2.410 ms
 7  beacon1.netnw.net.uk (194.66.26.42)  2.458 ms  2.561 ms  2.404 ms
```

Figure 13: July Traceroutes

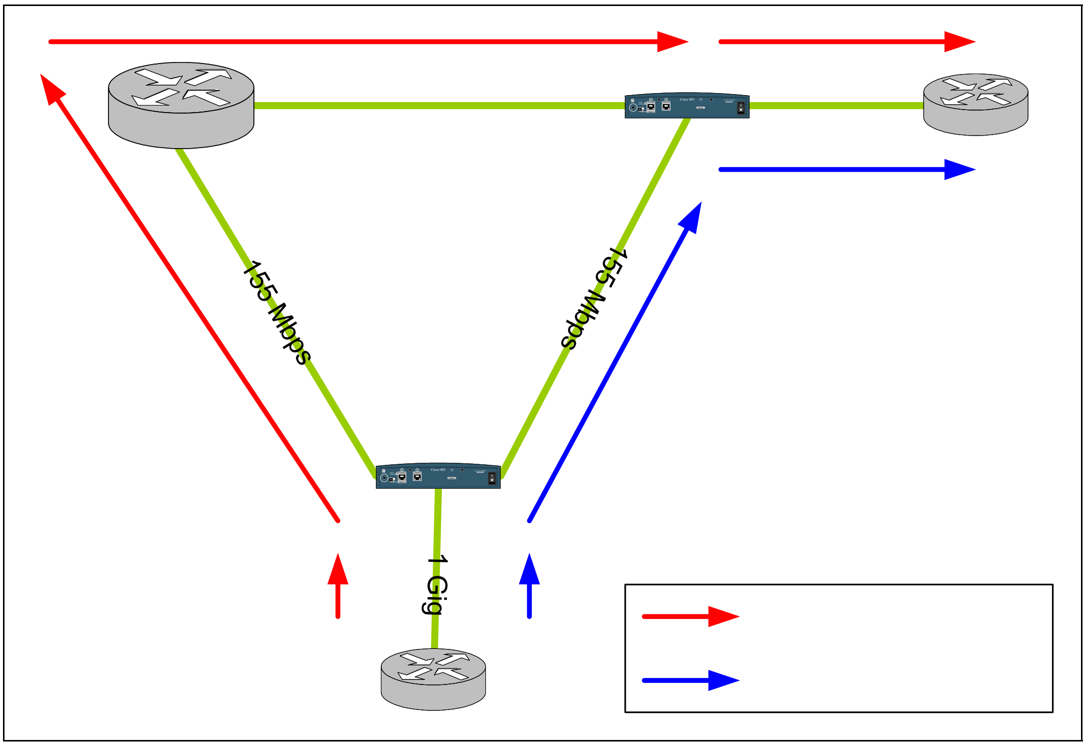

Figure 14: Paths Taken During Loading

7.3 Conclusions

Despite the changes to the network devices, the tests showed that the new devices (and their associated scripts) seemed to work well. The only confusion was with the rather odd routing reported in section 7.2, which we were unable to explain.

8 Trial Conclusions

The trials that we have completed as part of the QoS Development Project have proven extremely useful in terms of being able to understand the complexities and overall usefulness of supporting Quality of Service within the University environment.

Our main concluding points are as follows:

- When heavily loaded, VoIP calls were still possible. The loading on the network had to be significant in order to have major impact on the call. This highlights a number of points: 1) the VoIP phones handle abnormalities in the traffic flow well, 2) the requirements of a single VoIP flow in the context of the bandwidth of an access link are insignificant. This perhaps points to the use of a different measurement metric in the future

- The traffic used to load the network is highly significant. TCP's congestion control helps to keep some of the link bandwidth available, and as such does not really help here. When we used UDP, we have significantly "better" results in terms of being able to overload the link. In a number of instances, using UDP to load the link meant that calls were simply not possible

- On the whole, the equipment within our infrastructure survived, and did what was asked of it. The Cisco 805's were perhaps not up to the job of supporting the SAA measurements

- We were surprised at how difficult it was to gain access to the QoS statistics on different devices, and in some instances how crude the counters were, and how poor that granularity was

- Once the traffic that we were creating was outside of our management domain, things were very hard, if not impossible, to monitor. This brought home to us how difficult the overall end-to-end problem is

- There may be some rather interesting link sharing going on, but we were unable to prove this in any detail

There are many areas of future work that we would like to tackle. They include:

- **Video trials** – The more stringent requirements for video might enable us to see problems from an applications perspective, with less network loading. However, it is likely that (as with VoIP) multiple streams would be required to really have an impact

- **IPv6** – Given that JANET are now offering a full dual stack IPv4 / IPv6 service, it would be interested to look at the provision of EF and LBE for both IPv4 and IPv6

- **Access technologies / wireless** – many people in the networking and QoS community believe that the "QoS issue" is really an issue for the edge of the network, where both wired and wireless technologies with much smaller amounts of bandwidth are in use

- **Multidomain** – as was mentioned in the conclusions above, the multidomain issue is significant. In the context of the JANET QoS Development Project, the (small number) of partners were able to sit down together and make decisions about traffic levels, trust relationships, etc… without too much difficulty. In a real world scenario, policies and SLAs would be required in order to ensure fairness. Should the provision of these policies be dynamic in nature, this would imply a whole range of further issues

9 References

[1] The JANET QoS Technical Framework Document, available from http://www.ja.net/development/qos/TechnicalFramework.pdf

[2] Lancaster / C&NLMAN Deliverable 1: Initial Findings, available from http://www.ja.net/development/qos/qos_dev.html

[3] Net North West Partner Tests Document, available from http://www.ja.net/development/qos/qos_dev.html

10 Appendix 1: C&NLMAN Configuration

```
# Fire up the old diffserv
enable diffserv examination
enable diffserv examination ports all
# The mapping for DSCP values to hardware queues is as follows
# 00(000000)->07(000111) == QP1 - Normal
# 08(001000)->15(001111) == QP2 - LBE
# 16(010000)->23(010111) == QP3 - Normal
# 24(011000)->31(011111) == QP4 - Normal
# 32(100000)->39(100111) == QP5 - Normal
# 40(101000)->47(101111) == QP6 - IP Premium
# 48(110000)->55(110111) == QP7 - Normal
# 56(111000)->63(111111) == QP8 - Net control
#
# For each of the queues, set min and max bandwidths and priority
# 46 - IP Premium, RFC 2598
config qosprofile QP6 minbw 5% maxbw 5% priority MediumHi
# 08 - Less-than best efforts
config qosprofile QP2 minbw 1% maxbw 100% priority Low
# Make QP1 (Best-efforts) "more important" than QP2 (LBE)
config qosprofile QP1 minbw 0% maxbw 99% priority LowHi
# Make QP3,4,5,7 the same as QP1
config qosprofile QP3 minbw 0% maxbw 99% priority LowHi
config qosprofile QP4 minbw 0% maxbw 99% priority LowHi
config qosprofile QP5 minbw 0% maxbw 99% priority LowHi
config qosprofile QP7 minbw 0% maxbw 99% priority LowHi
# QP8 is network control. Allocate between 5 and 10% with maximum
# transmission priority...
config qosprofile QP8 minbw 5% maxbw 10% priority highHi
#
#
# Port control
# Turn on the BOFH ability to re-write the DCSP codepoint
enable dot1p replacement ports all
enable diffserv replacement ports all
# The ports on the black diamond are as follows…
# 1:1       10G Janet NIU
# 2:1       10G Lancaster University NIU
# 3:1       10G Carlisle
# 4:1       10G Daresbury NIU
# 8:1       1G CLEO
# 8:2       1G UKMS
# 8:3       1G 7206 Tunnel Endpoint
# 8:4       1G Lancaster University
# 8:5       1G 3550 Blackpool
# 8:6       1G St Martins
# 8:7       1G CEH Lancaster
# 8:8       1G Janet
# We need to re-mark any traffic not coming from either JANET or
# Lancaster University to DCSP 0
# Thanks to the way extreme's work, we have to first replace the
# dot1p value and then use that to drop the diffserv CP
config ports 3:1 qosprofile qp1
config ports 4:1 qosprofile qp1
config ports 8:1 - 8:3 qosprofile qp1
config ports 8:5 - 8:7 qosprofile qp1
# so then, if I've read this right, any traffic coming in on these
# ports should now have their dot1p value set to 0, we now need to
# set the diffserv CP to 0
config diffserv replacement priority 0 code_point 0 ports all
```

11 Appendix 2: Lancaster University BAR Configuration

```
! lu-bar QoS config
! There seems to be disagreement on whether you need to enable cef
! for Qos
! ip cef
!
! ports are as follows,
! GigabitEthernet0/0 - Uplink to JANET
! GigabitEthernet0/1 - Downlink to Campus Network
!
! Classify the traffic passing through the router
! AF - Assured forwarding, router updates traffic.
! IN - Inbound traffic from CANLMAN
! PREMIUM - Traffic allowed to mask DSCP values to 46
! REMARK - Catch anything else and remark its DSCP value
!
class-map match-all AF
      match ip dscp 48
      match access-group 199 in
class-map match-all IN
      match ip dscp 46
      match input-interface GigabitEthernet0/0
class-map match-any PREMIUM
      match ip dscp 46
      match access-group 190 in
class-map match-any  REMARK
      match input-interface GigabitEthernet0/1
!
! Create policies for each of the classes of traffic defined above.
! PREMIUM - gets 50Mbps (5% of 1gig)
! AF - gets 15% of remaining
! REMARK - sets the DSCP of any un-authorised traffic to 0
! DEFAULT - catches anything else and remarks to 0
policy-map LANCS
      class PREMIUM
            priority 50000
      class AF
            bandwidth percent 20
      class REMAP
            set ip dscp 0
      class DEFAULT
            set ip dscp 0
policy-map CANLMAN
      class IN
            priority 50000
!
! Then the policies need to be applied to the interfaces
! WARNING, don't just copy & paste this bit as it will break other
! interface settings
interface GigabitEthernet0/0
      service-policy output LANCS
interface GigabitEthernet0/1
      service-policy output CANLMAN
```

12 Appendix 3: Juniper Configuration

```
interfaces {
    /* Created from IOS Interface: gigabitethernet0/0 */
    ge-0/0/0 {
        description "Uplink to CANLMAN BAR";
        link-mode full-duplex;
        unit 0 {
            family inet {
                filter {
                    input 101;
                    output 102;
                }
                address 194.81.46.2/30;
            }
        }
    }
    /* Created from IOS Interface: gigabitethernet0/1 */
    ge-1/3/0 {
        description "Router network";
        unit 0 {
            family inet {
                filter {
                    input classify;
                }
                address 194.80.32.1/25;
            }
        }
    }

}
class-of-service {
    interfaces {
        ge-0/0/0 {
          scheduler-map LUNI-Scheduler;
            unit 0 {
                rewrite-rules {
                    dscp default;
                }
            }
        }
    }
    scheduler-maps {
        LUNI-Scheduler {
            forwarding-class network-control scheduler NC-scheduler;
            forwarding-class best-effort scheduler BE-scheduler;
            forwarding-class assured-forwarding scheduler AF-
scheduler;
            forwarding-class expedited-forwarding scheduler EF-
scheduler;
        }
    }
    schedulers {
        EF-scheduler {
            transmit-rate percent 10;
            priority high;
        }
        NC-scheduler {
            transmit-rate percent 5;
            priority low;
        }
```

```
        BE-scheduler {
            transmit-rate percent 80;
            priority low;
        }
        AF-scheduler {
            transmit-rate percent 5;
            priority low;
        }
    }
}
firewall {
    policer be-policer {
        if-exceeding {
            bandwidth-limit 100m;
            burst-size-limit 15k;
        }
        then loss-priority high;
    }

    filter 190 {
        /* access-list 190 permit ip 0.0.0.0 255.255.255.0 any
        access-list 190 permit ip host 148.88.3.10 any */
        term T1 {
            from {
                source-address {
                    0.0.0.0/0.0.0.255;
                    148.88.3.10/32;
                }
            }
            then accept;
        }
    }x
    filter 199 {
        /* access-list 199 permit ip host 194.80.32.1 any */
        term T1 {
            from {
                source-address {
                    194.80.32.1/32;
                }
            }
            then accept;
        }
    }
    filter classify {
        term expedited {
            from {
                source-address {
                    148.88.0.0/24;
                }
                dscp ef;
            }
            then {
                count expedited;
                forwarding-class expedited-forwarding;
            }
        }
        term expedited-2 {
            from {
                source-address {
                    148.88.0.60/32;
                    148.88.0.61/32;
```

```
                }
            }
            then {
                count expedited-2;
                forwarding-class expedited-forwarding;
            }
        }
        term confnat {
            from {
                source-address {
                    10.36.168.32/27;
                }
            }
            then {
                count confnat-policy;
                routing-instance confnat;
            }
        }
        term best-effort {
            then {
                count best-effort;
                forwarding-class best-effort;
            }
        }
    }
```

13 Appendix 4: Traffic Monitoring Source Code

14 Appendix 5: Remarking Voice Calls