The Use of QoS in the SSDN Interconnect and the Glow Project

Preamble

The Glow Project aims to deliver a rich teaching and learning environment with many tools and features over high-bandwidth links to Scottish Local Authorities for distribution to the Schools sector. The Glow project is currently under development and there is comparatively little "Glow" traffic as such, although a pilot involving 24 schools and 8 Local authorities has just been completed. In future Glow will provide streamed or prepositioned audio and video and implement net and videoconferencing and remote teaching as part of a wider package that also includes lower bandwidth applications.

The Scottish schools digital network (SSDN) Interconnect is the high bandwidth network over which Glow material will be delivered to the local authority networks. Circuits range from 45 to 1000Mbit/s. Currently this network is provided as an extension of the JANET network in Scotland which is administered by the United Kingdom Education and Research Networking Association (UKERNA). Learning and Teaching Scotland (LTS) is carrying out a re-provisioning exercise for the Interconnect. As part of this process LTS needs to consider bandwidth issues on the Interconnect.

JANET
JANET is the network dedicated to the needs of education and research in the UK. It connects the UK's education and research organisations to each other, as well as to the rest of the world through links to the global Internet. In addition, JANET includes a separate network that is available to the community for experimental activities in network development.

UKERNA
UKERNA manages the operation and development of JANET on behalf of JISC (Joint Information Systems Committee) for the UK Further and Higher Education Funding Councils. JISC also works in partnership with the Research Councils. UKERNA is government funded, with the primary aim of providing and developing a network infrastructure that meets the needs of the education and research communities.

QoS on the SSDN Interconnect Backbone

Traditionally Quality of Service (QoS) has provided a bandwidth solution by providing temporary channels of guaranteed bandwidth for marked priority traffic, with priority levels 0 (best effort) through 7 (network control reserved traffic). This often leads to packets with a lower priority being dropped although interactive voice (level 6) and interactive video (level 5) usually suffer minimum degradation.

UKERNA defines the classes slightly differently and has a less than best effort (LBE) classification.

With a view to deciding whether there is any advantage in implementing QoS is a valid solution to bandwidth shortage, LTS delegates (Jim Buchan and Ian McLean) have been attending meetings and participating in activities of the JANET QoS Phase 2 steering group.

QoS can be ineffective and even counter-productive on well-provisioned backbone networks, such as the SSDN Interconnect. ("Premium service on a well-provisioned network would do little to change packet forwarding under normal conditions" – Internet 2 QoS Working Group.) See, for example:

- http://www.oreillynet.com/pub/a/network/2002/06/11/platform.html
- http://qbone.internet2.edu/papers/non-architectural-problems.txt
- http://www.bricklin.com/qos.htm

QoS can provide advantages in networks that are under-provisioned (but not extremely under-provisioned). It provides little or no advantage on over-provisioned or well-provisioned networks. This was discussed at some length at the JANET QoS Phase 2 group face to face meeting on the 7th November 2006. Steve Williams presented an argument using Duncan Rogerson's slides which most effectively demonstrated that QoS would be an overhead on the JANET backbone that would provide little or no advantage.

"The idea to provide QoS only on under-provisioned links and hence to focus on end-site access links was approved" – Victor Olifer, JANET Phase 2 QoS Third Face to Face Meeting Minutes.

Traffic on the SSDN Interconnect

Measured over a typical 30 day period, the average and maximum traffic flows aggregated over all the sites on the entire SSDN Interconnect were:

Average from site:	37.54Mbit/s
Average to site:	82.96Mbit/s
Maximum from site:	330.08Mbit/s
Maximum to site:	984.65Mbit/s

It is not currently possible to analyse all of this traffic, which is coming from and going to 47 separate sites. However, it is reasonably safe to assume that only a very small percentage is shaped. Most of it is Internet traffic and QoS has not been generally deployed in the public Internet.

In practice, most modern routers are configured so that when a packet is forwarded from an interface with queuing, packets requiring low jitter - e.g. Voice over IP (VoIP) - are given priority over packets in other queues. Typically, some bandwidth is allocated by default to network control packets (e.g. ICMP and routing protocols), while best effort traffic might simply be given whatever bandwidth is left over. This is normal behaviour and is not regarded as QoS traffic shaping.

Also QoS requires an end-to-end agreement between all parties involved in the connection and needs to be supported by all routers on the route.

LTS can observe and analyse traffic coming from or going to the LTS offices at Dundee and Glasgow and the CDI traffic at Scolocate. The connection circuits to the Glasgow and Dundee sites are LES 100 whilst the Scolocate site is connected by means of a 1Gbit/s Ethernet provision on dark fibre. Traffic levels are as follows (all figures in Mbit/s):

	Average from site	Average to site	Maximum from site	Maximum to site
Learning and Teaching Scotland (Dundee)	0.28	0.31	10.37	16.19
Learning and Teaching Scotland (Glasgow)	1.27	4.41	11.85	55.53
Learning and Teaching Scotland (Scolocate)	4.98	0.83	56.62	11.28

No shaped QoS traffic was detected at any of the three sites. Even videoconferencing using JVCS generates best effort (free for all) traffic. As with the JANET network as a whole, there seems to be little or no advantage in implementing QoS as a backbone service on the SSDN Interconnect. If individual schools or Local Authorities (end-site access links) have a specific requirement this can be dealt with on a case-to-case basis.

There is a requirement that any end-site access QoS connection that is implemented can be made to work from end to end across the SSDN Interconnect community of sites. This is best achieved if any Scottish Local Authority that chooses to use QoS uses the same QoS scheme as any other Scottish Local Authority and the JANET core network passes packets without affecting any QoS marking.

QoS on Local Authority Networks

QoS is not seen by Local Authorities as a general answer to low-bandwidth problems. A school with a 128 Kbit ISDN connection is going to have problems with videoconferencing and rich media streaming no matter how much traffic shaping is applied. To date, bandwidth issues have arisen in Local Authority networks due to low-bandwidth ISDN and ASDL connections. It is estimated that as much as 30% of all schools in Scotland have links with 2 Mbit/s or less bandwidth.

Where a school has sufficient bandwidth (typically defined in the Glow Project as 4Mbps SDSL or better) QoS traffic shaping is typically used in multicast routing and VoIP traffic. The following is an extract from the configuration of a level three Catalyst 3550 switch in a primary school in a well-connected Authority. Descriptions have been altered to protect anonymity.

```
ip multicast-routing
mls qos map cos-dscp 0 8 16 26 32 46 48 56
mls qos min-reserve 5 170
mls qos min-reserve 6 10
mls qos min-reserve 7 65
mls qos min-reserve 8 26
```

```
mls qos
!
!
interface FastEthernet0/1
 description Link to Hub Secondary
 no switchport
 ip address 10.250.250.50 255.255.255.252
 ip pim dense-mode
 load-interval 30
 speed 10
 mls qos trust cos
 auto qos voip trust
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
!
!interface FastEthernet0/22
 description Voice Vlan
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport mode trunk
 switchport voice vlan 100
 no ip address
 shutdown
 mls qos trust device cisco-phone
 mls qos trust cos
 auto qos voip cisco-phone
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
 spanning-tree portfast
!
interface FastEthernet0/23
 description Voice Vlan
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport mode trunk
 switchport voice vlan 100
 no ip address
 mls qos trust device cisco-phone
 mls qos trust cos
 auto qos voip cisco-phone
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
```

```
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
 spanning-tree portfast
!
interface FastEthernet0/24
 description Voice Vlan
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport mode trunk
 switchport voice vlan 100
 no ip address
 mls qos trust device cisco-phone
 mls qos trust cos
 auto qos voip cisco-phone
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
 spanning-tree portfast
!
```

The following is an extract from the configuration of the level three Catalyst 3550 switch in the hub secondary school to which the above primary school connects:

```
ip multicast-routing
mls qos map cos-dscp 0 8 16 26 32 46 48 56
mls qos min-reserve 5 170
mls qos min-reserve 6 10
mls qos min-reserve 7 65
mls qos min-reserve 8 26
mls qos
!
!
interface FastEthernet0/3
 description Link to Primary School
 no switchport
 ip address 10.250.250.49 255.255.255.252
 ip pim dense-mode
 load-interval 30
 mls qos trust cos
 auto qos voip trust
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
!
interface FastEthernet0/22
 description Voice Vlan
```

```
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport mode trunk
 switchport voice vlan 100
 no ip address
 mls qos trust device cisco-phone
 mls qos trust cos
 auto qos voip cisco-phone
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
 spanning-tree portfast
!
interface FastEthernet0/23
 description Link to Secondary Power switch
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust cos
 auto qos voip trust
 wrr-queue bandwidth 20 1 80 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
```

Use of QoS between Local Authorities

I pursued this information by telephone and email enquiries. LTS has recently sent out a number of formal surveys and I was worried that another one might not be appreciated. As far as I can determine no QoS partnerships currently exist between Local Authorities, and QoS is used internally. This conclusion agrees with the lack of QoS traffic on the SSDN Interconnect backbone.

At the JANET QoS Videoconference meeting on Thursday, 22nd February it was suggested that Scottish Local Authorities be contacted so that two schools in separate Authorities could be nominated as test sites in a project to set up an inter-Authority QoS partnership. The following letter from Jim Buchan was sent to Local Authority technical contacts through email:

"QoS Tests Involving the SSDN Interconnect

Learning and Teaching Scotland (LTS) has been participating in the JANET Quality of Service Steering Group and is now keen to test the use of QoS over the SSDN Interconnect between two schools in different Local Authorities. This implies that QoS marked traffic will transit between the school via the SSDN Interconnect and the wide area networks of the participating LAs.

We need to identify two schools within Local Authorities that are willing to participate in such a test. Work would involve configuring a QoS link with the assistance of LTS and UKERNA staff. Benefits would include the ability to send high priority traffic such as video-conferencing or VoIP between the end points with minimum degradation. The ideal situation would be to identify two schools that are geographically distant but are cooperating on a joint educational project, but this is not essential. The schools in question would need to be connected to their Local Authority through 2Mbit/s Broadband links (or better). SDSL is preferred to ADSL for QoS links.

QoS identifies and marks priority traffic such as video and speech and guarantees a bandwidth allocation for that traffic. A QoS link is set up during (for example) a videoconferencing session and torn down when the session is finished, because unused bandwidth allocated to non-existent priority traffic means less bandwidth for normal traffic flows. QoS offers guaranteed bandwidth for traffic that would otherwise degrade significantly if packet loss occurred.

Setting up QoS requires additional configuration of QoS-capable routers (typically a few extra commands) and should not involve a lot of extra work. UKERNA in particular has considerable expertise in this area. Part of the project would be to determine how QoS can be implemented with the minimum amount of involvement by the participating Authority.

If you feel this could be of benefit within your Authority and you are interested in participating in a trial project, please contact Ian McLean at LTS (i.mclean@ltscotland.org.uk) in the first instance and identify possible participating sites.

Jim Buchan"

In spite of follow-ups by both phone and email, no definite school site has been proposed. One Authority expressed an interest but needed to know exactly how much work was involved. On the other hand I have received only one definite refusal. Currently many Authorities are undertaking major network redesign and reprovisioning exercises and are not prepared to commit themselves to any additional projects.

I believe that at least one school site will be identified before the end of this project. The ideal is to set up two sites in separate Authorities, but a QoS partnership between LTS and a single site is a fall-back position. However, I do not think that a QoS partnership will be in place across the SSDN interconnect before the deadline for this report.

Additional Solution to Bandwidth Problems - The Content Delivery Infrastructure

The Content Delivery Infrastructure (CDI) project, although a separate project in its own right, comes under the Glow umbrella. CDI addresses bandwidth issues through providing caching and prepositionioning, and through splitting high bandwidth streams at the LAN level—typically at edge content engines (CEs) in schools—rather than requiring multiple streams over the Internet or Local Authority internetwork. CDI is seen as complementary to QoS. The two do different things, but both in their own way address bandwidth issues.

LTS monitors CDI traffic and efficiency gains through the Content Delivery Managers (CDMs) which are installed at Scolocate, and through the bandwidth

statistics available at ssdn.netsite.ja.net. No serious congestion has been detected on the backbone network.

Implementation

The CDI structure uses the Cisco Application and Content Networking System to deliver strategic applications and content. Two Content Distribution Managers (CDMs) implemented on Cisco xxxx devices provide central content and device management for the system. The CDMs are located at a secure location at Scolocate, Edinburgh.

Hub and edge Content Engines (CEs) are located at Scolocate, at all Local Authority headquarters, and in every Local authority school in Scotland. In larger Local Authorities a root CE connects to hub CEs in secondary schools each of which in turn connect to a number of edge CEs in primary schools. In smaller Local Authorities edge CEs connect directly to the root CE. Root and hub devices are Cisco xxx CEs. Edge devices are Cisco 511 CEs. The devices are functionally identical except that xxx CEs have dual hard disks.

The CEs in every school serve client requests for content. Curriculum related content provided by establishments such as LTS and the Scottish Qualifications Agency (SQA) is held on Serendipity MYiNTERNET servers at Scolocate and at each Local Authority. This enables such material to be prepositioned in a CE at a school so that it is available on the school LAN on demand. Educational content can also be acquired from the Internet via the LTS root CEs at Scolocate, which act as global acquirers for both cacheable material and media streams. Figure 1 shows the CDI structure for a large Local Authority.

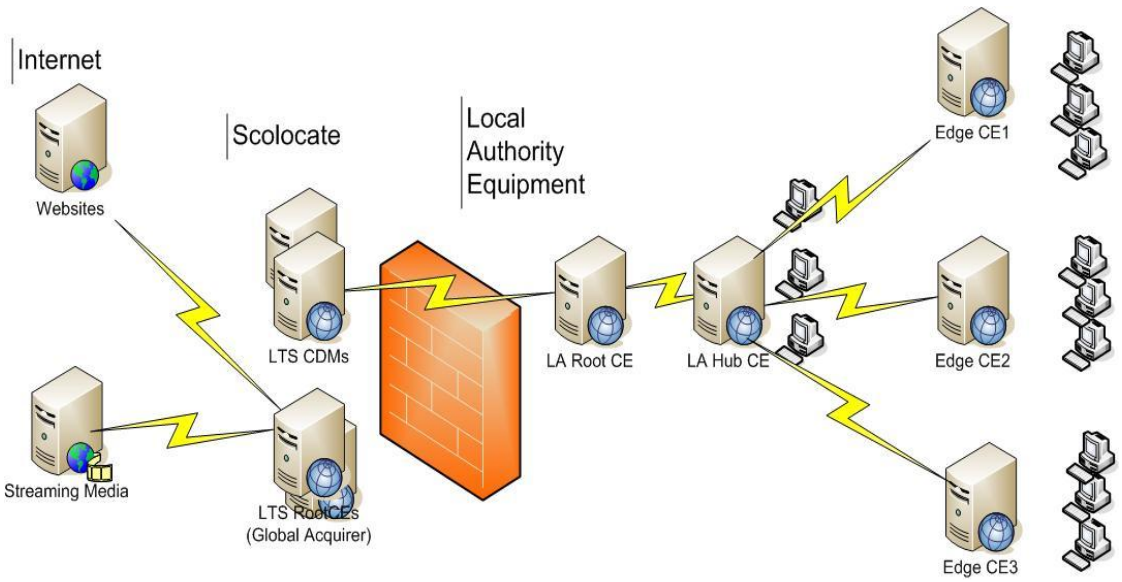

Figure 1 CDI Structure

Network Structure

Connectivity is achieved over the SSDN Interconnect, which is a public network, so all CEs have external IPv4 addresses in addition to internal IPv4 addresses on Local

Authority networks. Content can be distributed to a single specified CE, or to a number of CEs in a content channel. Each Local Authority network is autonomous.

The system is designed so that it requires the minimum number of ports to be opened in Local Authority firewalls. Port 443 is the only incoming port that is required so that the CDMs can manage the CEs. Ports need be opened only to the small number of devices (/28 network) at Scolocate. The system is a pull system. Local Authorities can pull information on request. Information is never pushed to Local Authority networks. Bandwidth Savings

The CDI system alleviates low-bandwidth restrictions within Local Authorities in a number of ways:

- Websites and website elements can be cached locally and subsequent request satisfied by cache. Although many modern websites are dynamic and cannot be cached in their entirety, elements such as website graphics can be cached. LTS monitors efficiency savings at the CDMs at Scolocate. Currently savings are typically upwards of 15%, which is significant but not revolutionary. However, as the CDI system is universally implemented in all Local Authorities this figure is expected to rise.
- Prepositioning guarantees that selected material is available locally on demand. Cached material might be overwritten but prepositioned material is guaranteed to be there. Material can be prepositioned several weeks in advance. Prepositioning is particularly significant for schools on very low bandwidth links because content-rich material can be prepositioned overnight, and even over several nights (unfortunately this does not work for schools that switch their ISDN connections off at 4:00pm).
- The CEs act as stream splitters for streamed media. Thus if twenty people in a school need to see a video or listen to a streamed audio transmission, only one stream need be accessed by the school. The CE then splits this into twenty live streams. If material is to be sent across a Local Authority only a single stream needs to be accessed by the Local authority Root CE, which can split it into a stream to each hub CE. The hub CE can then send a stream to each edge CE which splits this to individual PCs. The bandwidth savings are potentially enormous.

Logging and Filtering

The CDI system honours any content or address filtering in place within a Local Authority, although the Glow project may also provide a national filtering service. Transaction logging within the CEs gives a detailed description of all activity on a CDI network.

Conclusion

Arguably the best way to provide an acceptable service quality is through over-provisioning of bandwidth rather than traffic shaping, priority levels, Expedited Forwarding, Assured Forwarding or policing.

There appears to be little demand for traditional QoS techniques in high-bandwidth backbone networks such as the SSDN Interconnect. QoS is used typically within a Local Authority network but not between Local Authorities. Typically, QoS shaping is implemented if VoIP is used.

The CDI project provides a method of providing bandwidth savings and more efficient use of limited bandwidth. It is complimentary to QoS and is particularly relevant in subsidiary networks where low bandwidth causes problems. Efficient caching, prepositioning and streamed media splitting compliment QoS in that each addresses a different bandwidth issue.

If QoS is to be used between Local Authorities then a QoS partnership requires that both Authorities use the same markings. This agrees with the current philosophy that the backbone network should not re-mark QoS, but should be transparent.

Currently, Local Authorities have shown little interest in a cross-Authority QoS project. However, I shall continue to pursue this.

Ian McLean