Feasibility of building an all-IP network — the BBC NGN project

Martin Nicholson, Steve Westlake and Yuan-Xing Zheng
British Broadcasting Corporation

The BBC recently carried out a Next Generation Network (NGN) pilot to examine the feasibility of providing a single IP network for both broadcast and enterprise traffic.

Although the pilot was a success and a useful learning experience for the teams involved, they identified various hardware problems that need to be solved and the limitations of the control system. So, although the NGN pilot did not provide all the answers, it certainly was a good starting point.

An SDH (Synchronous Digital Hierarchy) network has been serving the BBC well for a number of years but support for the equipment is diminishing and the use of IP is becoming ubiquitous. This SDH network links all the major BBC sites and, being based on Raman amplification, is known as the Raman network. The functionality of Raman includes video and audio contribution and distribution, the enterprise IP network and various E1 circuits.

The Raman topology includes two diverse paths to each site, in order that audio and video circuits have 100% redundancy and are able to continue in the event of the loss of any single site. At the destination site there is a changeover device to switch between the two paths in the event of a failure. This is illustrated conceptually below.

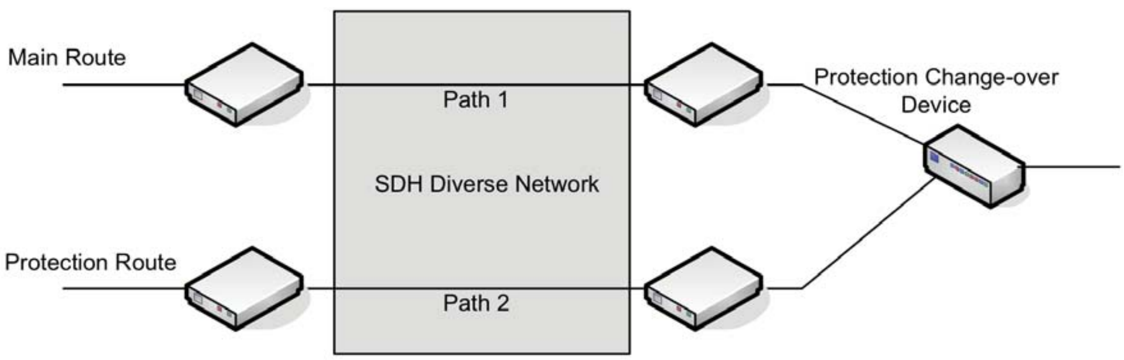

A High Availability Point to Point Circuit using SDH Technology

The NGN pilot aimed to prove that it is feasible to replace the Raman network with an IP network for both enterprise and broadcast traffic. To do this, a project was launched that involved multiple BBC divisions and several external project partners. The wide-area connectivity between sites was provided through Cable&Wireless Worldwide. The routers and other networking equipment were provided by Cisco whilst the IP video and IP audio codecs were provided by Cisco and Mayah respectively. The equipment to perform E1 emulation over IP was purchased from a company called Patapsco. All these devices were controlled by a Dimetis control system. Atos staff installed parts of the system and ensured that it was used for live traffic for the duration of the live phase of the pilot.

Challenges of IP

Internet Protocol (IP) is a widely-used general-purpose protocol used in many networks across the world and this helps to drive the costs down. IP however is not ideally suited for broadcast purposes and therefore there are many challenges that must be solved to enable the replacement of the existing SDH network. These challenges include:

- **Errors** – IP networks are typically associated with errors of various kinds where packets are lost or corrupted.
- **Latency / jitter** – IP networks by definition are packet-based with multiple traffic flows, meaning that buffers are required to eliminate the variation in packet arrival times (jitter) and this has the negative effect of increasing the latency (delays).
- **Dropped Packets / Quality of Service** – during periods of congestion, packets are dropped in a random fashion and, in the case of broadcast traffic, this is unacceptable. A well-defined Quality of Service policy is required to prevent this.
- **Volume of broadcast traffic** – typically IP is used in the IT world where streamed traffic flows such as VoIP are small relative to broadcast traffic. Uncompressed HD video for example needs around 1.5 Gbit/s.
- **Protection and self-healing** – some features useful for the IT industry are a hindrance for broadcasters. For example if a link fails then re-routing may occur and this could have the effect of congesting the remaining link(s), and all traffic flows could suffer rather than the few that would have otherwise suffered. Furthermore, the convergence time of protocols like OSPF is considerably longer than traditional methods such as SDI changeover relays.
- **Bandwidth-sharing between broadcast and enterprise** – the broadcast and enterprise traffic must be able to share the network in harmony without the enterprise traffic causing problems for the broadcast services. Likewise the broadcast traffic should not be able to increase beyond a defined rate such that the enterprise traffic always has the minimum bandwidth available.

Abbreviations

AES	Audio Engineering Society http://www.aes.org/
AF	Assured Forwarding
ASI	Asynchronous Serial Interface
CBWFQ	Class-Based Weighted Fair Queuing
DAB	Digital Audio Broadcasting (Eureka-147) http://www.worlddab.org/
E1	ITU telecom standard (European primary rate multiplex)
EF	Expedited Forwarding
FCS	Frame Check Sequence
FM	Frequency Modulation
IP	Internet Protocol
MAC	Media Access Control
MPLS	Multi-Protocol Label Switching
NGN	Next Generation Network
NICAM	Near-Instantaneous Companding And Multiplexing
OSPF	Open Shortest Path First
P2MP	Point-to-MultiPoint
PCM	Pulse Code Modulation
PE	Provider Edge
PHB	Per-Hop Behaviour
PSNR	Peak Signal-to-Noise Ratio
QoS	Quality of Service
RIP	Raman Interconnect Point
SDH	Synchronous Digital Hierarchy
SDI	Serial Digital Interface
THD+N	Total Harmonic Distortion plus Noise
TE	Traffic Engineering
ToS	Type of Service
VoIP	Voice-over-IP

The solution: Multi-Protocol Label Switching

In order to address these challenges, Multi-Protocol Label Switching (MPLS) was used. Typically, IP networks are treated as a "cloud" with undefined and varying routes through the network. By contrast, MPLS allows fixed routes, or tunnels, through the network, making the network more deterministic. Fixed paths improve the deterministic nature of the network by creating known routes with reasonably constant latency, whilst removing the possibility of undesired re-routing effects. In the pilot, one broadcast circuit was used per tunnel (but this does not have to be the case).

Defined routes through the network allow a similar approach to the current SDH network, where two diverse paths are used with a changeover device at the destination. It is possible to keep the diverse routes by replacing the existing SDH network with an MPLS network as illustrated below.

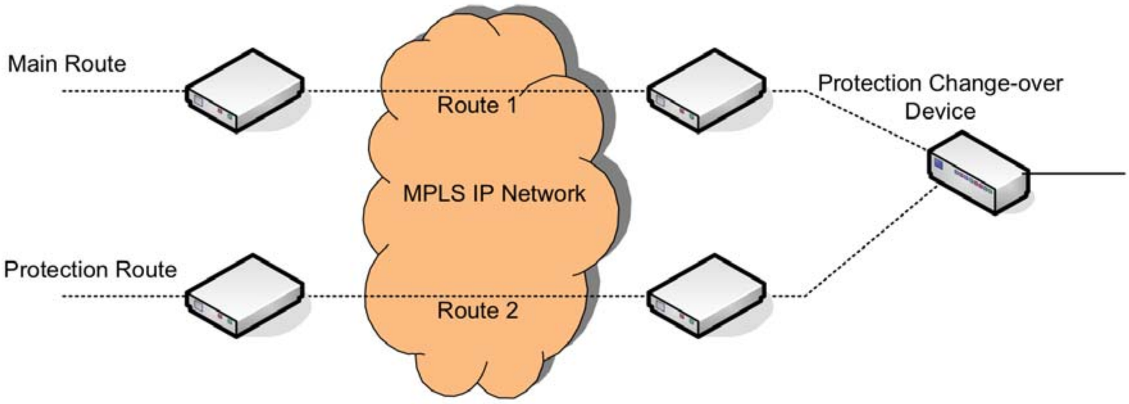

A High Availability Point to Point Circuit emulated over MPLS Technology

MPLS on its own is not enough to guarantee performance, so Traffic Engineering (TE) with point-to-multipoint (P2MP) routing was used to allow defined bandwidths on routes through the network. This, when combined with admission control for the available bandwidth, meant that it was not possible for the broadcast traffic to be oversubscribed and for congestion to occur. It should be noted that if traffic congestion is allowed to occur, then all streams in the same or lower priority class will suffer, not just the latest to join. The TE solution employed allowed tunnels to be dynamically joined or unjoined as required.

Point-to-multipoint routing, as the name suggests, enabled the possibility of having more than one receiver for each stream. It is important to note that this was not done using multicast, where receivers can dynamically join and leave because a more controlled method was required. P2MP allowed the control system to modify tunnels to include additional destinations without interruption to the existing stream. These tunnels remained active until they were stopped using the control system. This allows the control system to track exactly how much bandwidth is available on each link.

Drivers for a pilot

The pilot was necessary to confirm the theory and prove that an MPLS network is capable of replacing SDH. Three sites were chosen for the pilot, each site with two diverse paths. There were three key questions to answer:

- Can we carry broadcast traffic reliably over IP to match the current SDH platform?
- Can we do "anywhere-to-anywhere" routing in the IP domain?
- Can we successfully mix enterprise and broadcast traffic in a single network?

Network topology

The pilot included London Broadcasting House, London Television Centre and BBC Manchester, each with two diverse paths and two identical sets of equipment. The project temporarily installed

MPLS routers, video codecs, audio codecs, layer 2 switches and other equipment in the BBC sites. In addition, an MPLS router was installed in two Telco hub sites known as Raman Interconnect Point (RIP) North and South. These three BBC sites had 10 Gbit/s wide-area connectivity, provided to both RIP North and RIP South. The following illustration shows the conceptual topology.

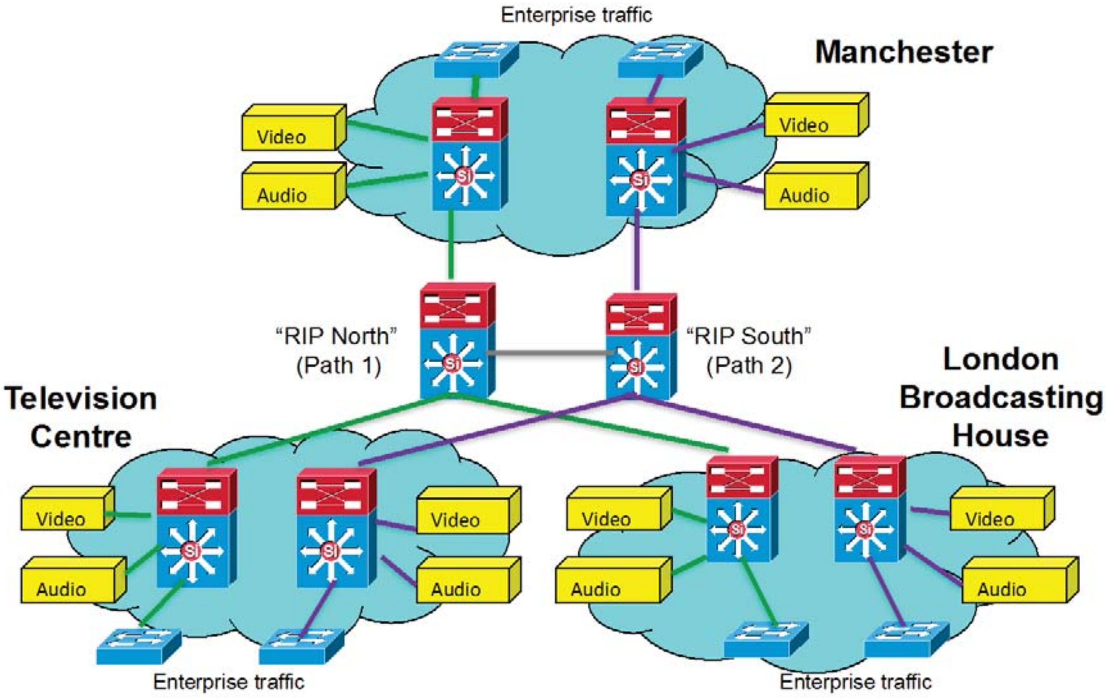

Major stages in the pilot

After the feasibility study and the partner selection, the pilot itself lasted from November 2010 until March 2011.

All the chosen equipment was delivered to the lab in Reading, west of London, where the team spent around two months in the lab-testing phase. The proposed initial configuration was applied and thorough testing of all the equipment was undertaken, including tests at network level, video level and audio level. It was in this phase that testing took place for different Quality of Service policies and one of them was chosen. Having all the equipment located in a single lab made problem-solving and configuration changes possible over short periods of time without the requirement of a large team. The lab-testing phase was where several precise measurements were made concerning the recovery time involved in the event of various deliberately-introduced faults.

The field-testing phase also lasted for around two months and featured different challenges. First and foremost was to check that all equipment functioned as expected after being moved from the lab to the correct site. The increase in latency due to the increased physical distance was measured, and practical issues such as different word clock sources for audio were solved.

The live phase lasted for approximately three weeks and saw the pilot network used for real audio and video traffic before being broadcast. The services included in the live phase consisted of permanent tunnels as well as "anywhere-to-anywhere" dynamically joined / left tunnels for audio and video contribution. The following services were included:

- The BBC 1 sustaining feed to Manchester used for the distribution of BBC 1 North West for all platforms.
- The NICAM radio distribution feeds used for the main FM radio channels to the North West and feeds to a range of DAB and Radio 5 transmitters in the North of England. It also included a Radio 5 sustaining feed for local radio stations in the area.
- Video and Music contribution circuits were available for bookings between the pilot sites.

Connecting the BBC enterprise network to the pilot was considered unfeasible so, during the live phase, enterprise traffic was simulated instead. This simulated traffic was designed to represent the protocols and the “bursty” nature of real traffic. As such, there was between 1 and 2 Gbit/s present constantly in the background in order to prove that it did not affect the higher priority broadcast traffic.

Codec testing

The codecs used for SD or HD SDI video over IP allowed for uncompressed video or for JPEG2k compression. The video latency, with no physical distance between the encoder and decoder, measured in the lab phase was:

- Uncompressed HD = 9.1ms ±0%
- Uncompressed SD = 9.8ms ±0.3%
- 150 Mbit/s JPEG2k = 150.61ms ±3%

There was an updated beta firmware available after the pilot for which the video latency was measured to be:

- Uncompressed HD = 10.05ms ±1%
- 150 Mbit/s JPEG2k = 95.96ms ±1%

As part of the video codec testing, the Peak Signal to Noise Ratio (PSNR) was measured for various bitrates when using JPEG2k compression for the transport of SDI video. This was undertaken in order to help decide which bitrate should be chosen for JPEG2k compression. The content used to measure the PSNR was the EBU HD test sequence because it featured a variety of scenes with different compression challenges. The following graphs show the PSNR of a sample from each scene of the test sequence against increasing bitrate for both luminance and chrominance. Interestingly a plateau is reached for the luminance *(bottom of this page)* at around 250 Mbit/s for luminance but chrominance *(top of next page)* does not reach a plateau, even by 700 Mbit/s.

Testing was also undertaken for ASI-over-IP codecs. The BBC currently has SDI-to-ASI codecs that use JPEG2k compression, meaning that ASI-to-IP codecs can be used with this existing infrastruc-

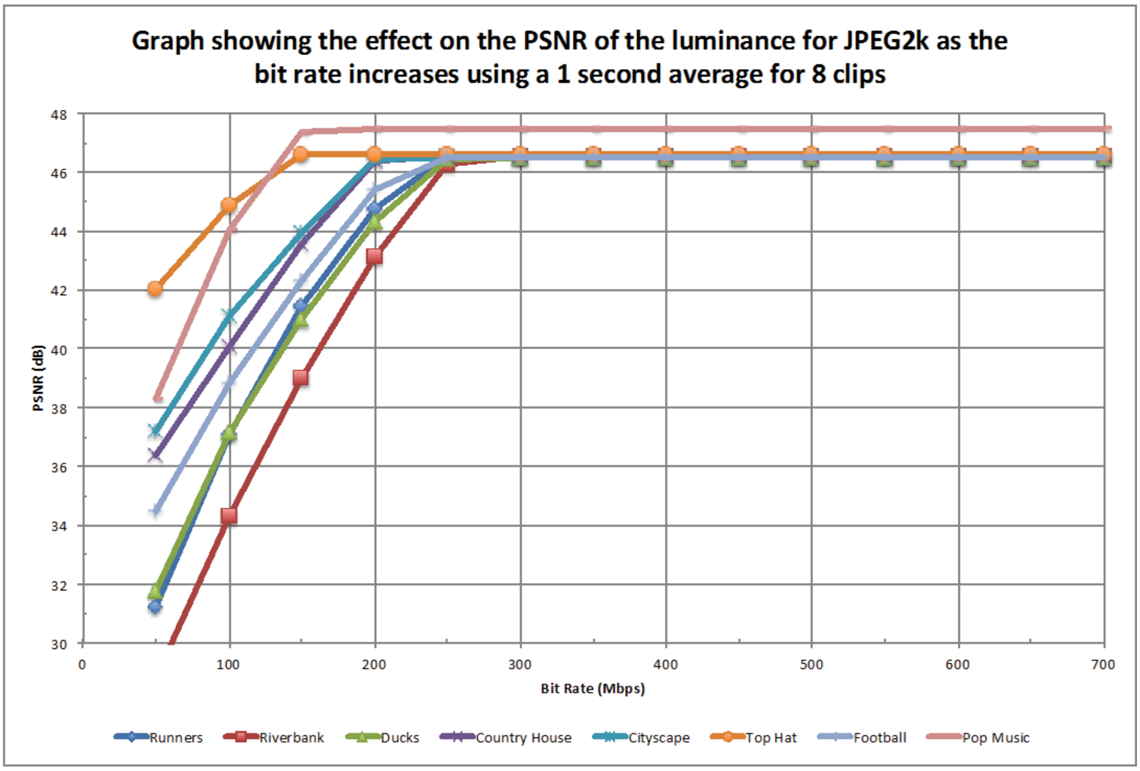

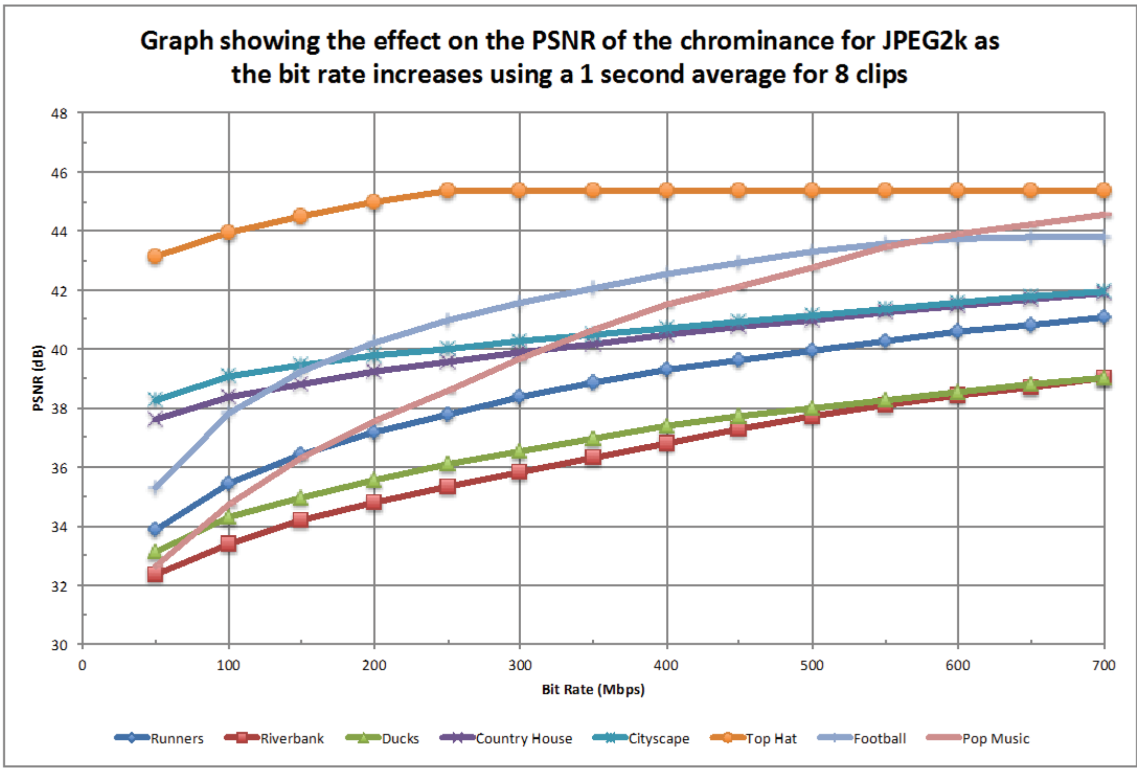

ture. The ASI codec latencies were measured in the lab phase and had no variation. The measured latencies were as follows:

- 35 Mbit/s = 12.11 ms
- 50 Mbit/s = 12.19 ms
- 90 Mbit/s = 12.09 ms

In addition to video codecs, there is a need for audio codecs. The pilot used AES3 audio-over-IP codecs, which were able to stream linear audio or use various compression algorithms. The latency of the audio streams varied depending upon which algorithm was in use. A variety of algorithms were measured in the lab phase and the results were as follows:

- AES3 transparent = 11.79 ms
- 24-bit linear PCM = 9.98 ms
- 16-bit linear PCM = 9.98 ms
- apt-X stereo = 13.02 ms
- G.722 = 36.04 ms
- G.711 = 34.35 ms ±2%

Due to the sensitive nature of AES3 audio for clocking purposes, extensive testing was undertaken on the audio output to check for audio consistency. To do this, a Prism Sound dScope was used in two modes: the “Channel Check” mode and the “Glitch Test” mode. The “Channel Check” mode checks that the audio path is transparent and, as such, only works with the AES3 transparent algorithm whereas the “Glitch Test” provides the Total Harmonic Distortion plus Noise (THD+N) figure for the audio path. The “Glitch Test” is therefore suited to measuring any algorithm and consequently was used most frequently. Using the dScope proved that the audio was not creating audible errors such as “pops” and “clicks” (often known as glitches). Imperfect clocking often causes glitches in AES3 audio.

The BBC uses NICAM-over-E1 tributaries for radio distribution, meaning that E1 emulation is required on the IP network to operate with existing equipment. E1 circuit emulation boxes were installed between London Broadcasting House and Manchester on both paths. These boxes encapsulated four E1 signals into a single 10 Mbit/s IP stream and, as expected, the boxes used adaptive clock recovery. This clock recovery featured a long time constant to adjust to changes, which resulted in a stable low-noise clock but caused problems which, although they did not greatly affect the pilot, were noted for future procurement. A clock change of 10 ppm could take ten minutes before the decoder output adjusted to the new frequency and in the case of a change of 20 ppm the decoder may not track without an interruption in the stream. This meant that an E1 circuit with more than a single hop was in danger of the buffers under-running. The pilot was therefore limited to single hops for E1 emulation in order to meet the BBC requirement to switch source encoders between the main and reserve chains. As with other equipment, this was something that the manufacturer fixed after the pilot, by providing a revised clocking algorithm.

Quality of Service

In an IP network, not all packets have to be treated equally. It is possible to assign *high priority* for the important traffic and *low priority* for the less important traffic. This technique is known as Quality of Service (QoS). In the pilot, the highest QoS priorities were assigned to Voice-over-IP (VoIP) and broadcast traffic while the lowest QoS priorities were for the enterprise traffic.

If there was ample bandwidth then, theoretically, QoS is not necessary because the network would not be congested. In reality, however, even with good network management systems in place, the utilisation measured on network links are averaged and as such these statistics will not pick up microbursts (short bursts which may hit peak utilisation). Microbursts are common where streams of bursty traffic use a single link. QoS is required to experience consistent end-to-end treatment of traffic types, even in times of congestion or microbursts.

The seven categories listed below were implemented by setting each device (IP adaptor) manually such that it included a Type of Service (ToS) number in the IP header of every packet sent. The exception to this was the management traffic, which used a separate ingress and egress layer-2 switch, so that it became the responsibility of the PE routers to set the ToS value. The seven QoS categories were:

- QoS 6: used for network control traffic;
- QoS 5: the lowest latency class of traffic used by the broadcast audio and the VoIP;
- QoS 4: the second lowest latency class used by the broadcast video and ASI;
- QoS 3: the highest enterprise level known as “critical” traffic used for enterprise video;
- QoS 2: the class for management traffic;
- QoS 1: an enterprise class used for the more important best-effort traffic;
- QoS 0: an enterprise class used for the less important best-effort traffic.

It is important to police or administer the admission control on the higher QoS classes or they can disrupt all traffic on the network. There were two main types of traffic used in the pilot: *Expedited Forwarding* (EF) and *Assured Forwarding* (AF).

RFC2598 defines the Expedited Forwarding (EF) Per Hop Behaviour (PHB). This class exhibits low latency and low loss characteristics. It is essentially a priority class and as soon as traffic is detected in this class, it will skip queues of traffic in any other classes.

RFC2597 defines the Assured Forwarding (AF) Per Hop Behaviour (PHB). This class is able to offer different levels of forwarding assurances for IP packets. It guarantees a certain amount of bandwidth to a class and allows access to extra bandwidth if available.

QoS policies assign each QoS level to a PHB type and the BBC pilot tested two QoS policies. The first was a CBWFQ (Class-Based Weighted Fair Queuing) egress policy. This policy enabled one class of Expedited Forwarding. *Table 1* shows conceptually how each of the seven QoS categories used was treated.

Table 1

Expedited Forwarding	Assured Forwarding			
QoS 5	QoS 4	QoS 3	QoS 2/QoS6	QoS 1/QoS 0
10% allocated	90% remaining	8% remaining	1% remaining	1% remaining

The second policy was a dual priority EF queue. This enabled two levels of Expedited Forwarding in addition to the above, as illustrated conceptually in *Table 2*.

Table 2

Expedited Forwarding		Assured Forwarding		
QoS 5	QoS 4	QoS 3	QoS 2/QoS 6	QoS 1/QoS 0
10% allocated	80% allocated	90% remaining	1% remaining	1% remaining

The first test found that congesting QoS 5 had the desired affect on both policies in that QoS 1 and QoS 0 suffered packet loss and increased latency. This is desired because QoS 5 is higher priority and represents the broadcast traffic whereas QoS 1 and QoS 0 are lower priority and represent enterprise traffic.

The second test found that congesting QoS 3 had the same effect as congesting QoS 5 in that QoS 1 and QoS 0 suffered packet loss and increased latency. This was the case for both CBWFQ and dual-priority EF policies. However a difference between the policies was identified during this test, namely that the CBWFQ policy suffered from increased latency in the QoS 4 class. This meant that using CBWFQ caused congestion in QoS 3 to affect the higher priority QoS 4 traffic. The latency increase was a sudden shift by approximately 60 ms that was severe enough to be noticed by viewers.

Clearly the second test proved that CBWFQ is not suitable for the pilot so the dual-priority EF policy was chosen instead. It is worth noting that not all router manufacturers support dual-priority EF queuing.

Packet sizes

Allocating bitrates to each QoS class is not as trivial as might be thought because the size of a packet physically “on the wire” and the size, as counted by routers, are different. The routers do not count all aspects of an Ethernet frame; for example they do not count the preamble or the inter-frame gap. The result of this difference can be significant and becomes especially noticeable when there are a large number of small packets. A closer look at an Ethernet frame *(Table 3)* illustrates which fields are counted and which are not.

To illustrate the extent of the difference between using large packets and using small packets, two example streams have been calculated using the same notional bitrate (as set on the device). The different sized packets produce a vastly different final “on the wire” bitrate. For example, a 500 Mbit/s stream using 64-byte packets has a final “on the wire” rate of 750 Mbit/s whereas the same stream using 1,386 byte packets has a final “on the wire” rate of around 512 Mbit/s. These examples are shown in *Table 4*.

Table 3

	Preamble	Start of frame delimiter	MAC destina-tion	MAC source	Ethertype or length	MPLS Label	Payload (data)	Frame Check Sequence (FCS)	Inter-frame gap
Size (bytes)	7	1	6	6	2	4 or 8	46-1500	4	12
On the Wire									
Routers									

Table 4

Bitrate set within the test device (bps)	Packet Size (Bytes)			Final "on the wire" bitrate (bps)
	As sent by the test device	Added overheads (+24 bytes)	Added MPLS label (+8 bytes)	
500,000,000	64	88	96	750,000,000
500,000,000	1,386	1,410	1,418	511,544,012

This is because, although a packet starts at a defined size, when transmitted 24 bytes need adding for the missing overheads: (i) the preamble; (ii) the start of frame delimiter; (iii) the Frame Check Sequence and (iv) the inter-frame gap. Furthermore, at the point of sending, the frame contains a native IP payload with no MPLS label; when the frame reaches the PE node, a label of 8 bytes is added. After these have both been added, the frame becomes 32 bytes longer.

This difference makes allocating bitrates to each QoS class difficult because knowledge of the packet sizes is required. When trouble-shooting, care should be taken in diagnosing the situation because the actual on-the-wire stream rates can be larger than reported.

Fixed routes versus anywhere-to-anywhere

In a static network (a network consisting of only fixed broadcast links), traffic management is implemented by planning and configuration control alone. The BBC NGN pilot used several instances of static MPLS tunnels configured outside of the control system. These tunnels used a fixed route through the network and can be thought of as the equivalent of traditional fixed broadcast links.

The pilot also used "anywhere-to-anywhere" routing where the control system manages the bandwidth and configures the devices on the network in order to create circuits. The device configuration includes the tunnel through the network and the codecs at each end. This responsibility makes the control system a tremendously important part of the network. Such an important system needs to be reliable but problems were found in the pilot that would need to be solved before final deployment.

For example, the failure of a codec was a problem because dynamically-configured circuits using that device could not be processed cleanly (for example, creating / modifying / deleting). The reason for this was that being unreachable meant that part of the way through the operation, errors would be reported on the control system and the process would abort. If attempts were made to process a circuit using devices that were not reachable, the circuit could end up somewhere between fully removed and fully working. In this situation, extensive manual intervention was required to restore all the devices back to the correct state.

The recommended action was to replace the faulty device in the network and then delete all the connections that were present on that device using the control system. This ensured the device was reachable despite not having any of the connections in memory, due to it being a replacement device. Being reachable meant that the control system was able to delete the entire connection as normal. The connections could then be manually re-established afterwards using the usual procedure on the control system. Clearly this recommended action was only possible if a device had already been identified as failed and was replaced before any configuration of that device was attempted.

The time taken to create, modify or delete a circuit was around 10 seconds, which the operators of the system found frustrating as they are used to traditional broadcast routers where switching happens in a few micro seconds.

Findings and conclusions

The pilot found that the MPLS network functioned as a valid alternative to the current SDH network with similar reliability. Permanent circuits could be configured with high availability and could be kept separate from the control system and could be managed manually. The desired anywhere-to-anywhere operation was feasible using the control system, and circuits could be set up in around 10 seconds. Enterprise traffic could be mixed with broadcast traffic on the same network with the desired QoS behaviour during periods of congestion. One key finding was that the broadcast traffic had to be in the Expedited Forwarding (EF) queue to guarantee latency during network congestion, meaning that the dual-priority EF queue was a more appropriate QoS policy.

As expected, the broadcast signals had higher latency compared to the current SDH network but could be used with few changes to operational practice. The time involved in creating anywhere-to-anywhere circuits would most likely lead to a mixed economy of permanent and anywhere-to-anywhere circuits.

Operational support is one area that would be likely to change in a fully-deployed NGN because a different skill set would be required compared with the current network.

Here are some answers to the questions raised earlier in this article, that the NGN pilot set out to answer:

Can we carry broadcast traffic reliably over IP to match the current SDH platform?

The answer to this question is “yes” but there are several caveats:

- The latency of broadcast signals is higher compared to the current SDH network. In some cases this will require some changes in operational practice but in most cases this should be of limited impact.
- The IP network upgrade to support MPLS at the many BBC sites will require significant upgrades on the currently deployed edge and core routers.
- The BBC network currently supports an extremely large number of broadcast circuits. The pilot stressed the MPLS system but was unable to create numbers even close to the estimated 8000 circuits, so there is some uncertainty on how such a big system will behave.
- It is expected that the IP network management teams will make the transitional leap required to support broadcast traffic.
- Some of the codecs used in the pilot need to show significant improvements in order to be deployed for service.

Can we do “anywhere-to-anywhere” routing in the IP domain?

The answer to the second question is also “yes” but again with several caveats:

- All the caveats to question 1 also apply for this second question.

Martin Nicholson is a Technologist at BBC Research and Development. He has undertaken a variety of projects whilst at the BBC, ranging from making television more accessible for the blind and partially sighted to innovative sports graphics for the 2012 London Olympics. Most recently he has been working on a next-generation network for the BBC and was heavily involved in a live pilot demonstrating one implementation of such a network.

Steve Westlake is the BBC Network Architect responsible for inter-site connectivity to support broadcast services. He first joined the BBC in 1971 and has been involved in several generations of network refresh, with the most recent pilot looking at the implications of broadcast services carriage over an MPLS infrastructure.

Yuan-Xing Zheng is a Lead Technologist in the BBC Network Architecture team. Her research interests include enterprise network design, video and audio over IP, IP monitoring & measurement and wireless communications. She is currently chairing the EBU Service Level Agreement (SLA) group, and is a member of the EBU Video over IP group (VCIP).

- The time taken to make or break an "anywhere-to-anywhere" connection in the IP domain is significantly higher than with a traditional video or audio matrix. This makes it suitable for routing occasional-use signals where time constrains are not as critical. It will be less suitable where the bookings between two sites tend to be high volume and time critical, such as feeds to network playout. It is highly likely that any practical deployment would feature a mix of "anywhere-to-anywhere" along with permanent circuits.
- The control system will become critical in managing the network and further development is required to reach a deployable system. Within the pilot there were only a small number of signals routed, yet there was evidence that the more the system was used, the more background maintenance was required to clean incomplete transactions. These incomplete transactions occasionally blocked sources and destinations, and could consume bandwidth.
- It should be noted that the control system manages the bandwidth and provides admission control for the anywhere-to-anywhere circuits. This is very different to a "plug and play" service that some users might dream about. If there were "plug and play" traffic then, realistically, this would not be allowed in the high-priority categories and therefore would fit within the lower priority categories, probably the best-effort category, like it does today.

Can we successfully mix enterprise and broadcast traffic in a single network?

The answer to the third question is a "yes" and the enterprise tests concluded that:

- We can achieve desired QoS protection and segregation between broadcast and enterprise traffic. This enables the ability to share the bandwidth between broadcast and enterprise uses to give greater flexibility.
- Different traffic priorities are possible within the enterprise flows for different services.
- Greater segregation of networks is possible allowing new services such as extranets.

To conclude, the BBC Next Generation Network pilot was a success and a useful learning experience for the teams involved. The teams identified the hardware problems that need to be solved, the limitations of the control system and were able to experiment with different network configurations. The NGN pilot did not provide all the answers, but it certainly was a good starting point.

This version: 31 March 2012

Published by the European Broadcasting Union, Geneva, Switzerland

ISSN: 1609-1469

Editeur Responsable: Lieven Vermaele

Editor: Mike Meyer

E-mail: tech@ebu.ch

The responsibility for views expressed in this article rests solely with the authors