JANET Content Delivery Infrastructure Trial Report

September 2003

Authors

Roger Bolam, Alex Roberts, Paul Wakefield, Steve Williams

JANET Content Delivery Infrastructure Trial Report

Contents Page

Preface

This report is based upon the findings of the JANET Content Delivery Infrastructure (CDI) Trial. The CDI trial was undertaken to prove the feasibility of deploying content delivery architecture on JANET. This report will be followed by a further report, that will recommend a Content Delivery Architecture for JANET.

1. Executive Summary

Following the conclusion of the Content Delivery Infrastructure (CDI) for SuperJANET: Requirements Study conducted by ESYS, the JANET CDI Trial was set up to follow up the recommendations of that report and to explore a delivery architecture that would scale to support the needs of the JANET community over the coming years.

In order to scale to the size required by the JANET community, the solutions considered are those that will support a hierarchical delivery architecture; consisting of a core server, or server farm, and a cache at organisation level that would then deliver streams locally to client media players. An architecture that did not cache content at the edge/organisational level would not be scalable on JANET due to the bandwidth demands of high quality video content.

Following a Call for Participation in the fourth quarter of 2002, four Higher Education (HE) and two Further Education (FE) sites were selected to participate in the CDI trial, with the University of Wales Swansea as the lead site. The trial commenced in January 2003.

As recommended by the ESYS report, Kasenna™ Mediabase software running on Silicon Graphics Inc®/IRIX (a version of UNIX for Silicon Graphics machines) and additional Intel®/Linux® platforms were deployed at the six trial sites and an SGI®/IRIX server was installed in the Reading Core Point of Presence (C-PoP).

In order to monitor the impact of streamed media on JANET, a monitoring infrastructure was installed in parallel. This was based on Cisco®'s Service Assurance Agent running on dedicated routers at the Reading C-PoP and at all the participating sites. Measurement data was gathered on a Structured Query Language (SQL) server for later analysis.

Content authorised through the British Universities Film and Video Council (BUFVC) / Managing Agent Advisory Service (MAAS) was made available to participating sites through a secured mesh of links. Sites outside the trial were not allowed access to the trial content or infrastructure.

The ESYS Report specifically recommended Kasenna™ Mediabase Network Edition as it supported most common delivery protocols in use (Windows®Media, Quicktime®, RealNetworks® and Moving Picture Expert

Group (MPEG)). However this software ceased development between the publishing of the ESYS report and the start of the JANET CDI Trial.

Following discussion between the United Kingdom Education and Research Networking Association (UKERNA), the BUFVC and the lead site, Kasenna™ Mediabase XMP SE was selected as the replacement software version. However, XMP SE only provides standards-based MPEG support. It was agreed that BUFVC/MAAS would provide MPEG4 encoded content for the trial.

During testing it became apparent that Kasenna™ Mediabase XMP SE was unable to cache MPEG 4 content. Due to a software bug in the current version, whilst XMP SE could serve MPEG 4 content, it was unable to store an incoming MPEG4 stream for local playback. Testing of XMP SE was suspended.

During the early stages of the trial the lead site had been investigating whether other European National Research and Education Networks (NRENs) were considering implementing similar architectures. Whilst there is considerable interest, most other NRENs are primarily looking at live streaming infrastructure and the Trans-European Research and Education Network Association (TERENA) had formed a Task Force, TF-Netcast, to investigate this area. Those NRENs that had considered video-on-demand provision were using a combination of Kasenna™, RealNetworks® and Microsoft® solutions.

Following discussion between UKERNA, the BUFVC and the lead site, a decision was taken to approach RealNetworks® UK, as the Helix™ range of products was seen to provide delivery architectures similar to those required. RealNetworks® agreed to work with UKERNA and provided licences to enable trials to take place. A Helix™ Universal Internet Server was installed at the Reading C-PoP on an Intel®/Linux® platform, and a Helix™ Universal Gateway was installed at all participating sites. Additional hardware was purchased for those sites that had previously only had SGI®/IRIX platforms, as Helix™ Universal Gateway was not available for these platforms.

Testing continued using the Helix™ platform with QuickTime® encoded content, and participating sites ran scheduled tests with single and multiple streams being viewed. Most tests consisted of sites watching between 10 and 30 simultaneous streams, but other load tests were done with sites watching over 100 simultaneous streams successfully across 100Mbit/s links.

The effect on the network of between 10-30 streams was deemed to be negligible at sites with a JANET connection of 100Mbit/s or more. The two FE colleges both had 2Mbit/s connections, and as expected these sites relied heavily on local caching to be able to successfully deliver media.

The architecture model investigated performed admirably. Content was delivered from local caches, so that JANET access links for the high speed connected sites saw only a negligible increase in load. For these higher

speed connected sites, the value of the cache to reduce access-link bandwidth was evident, and would become essential as usage of the system increased.

For lower speed (2Mbit/s) connected trial sites, once content was held in the local cache, sites could simultaneously watch large numbers of streams with a negligible increase in access-link traffic. Some traffic is required between the origin server and cache, even with fully cached material, as the cache examines the original content source during playback to ensure that a new version has not been uploaded, or that it has not been removed. If the original material has been removed, the cache will refuse access requests to playback that content – for rights' management reasons.

Investigation is continuing as part of an extension of the CDI Trial, funded by UKERNA, into certain key areas that it feels require a more in depth study before it would be possible to move forward.

Some of the areas undergoing further exploration are those in the recommendations section, but in particular the issue of a non-proprietary solution to a caching infrastructure is being investigated. Unlike the current situation with web HyperText Transfer Protocol (HTTP) caching where there is a choice between commercially licensed and free open-source solutions, stream caching products are normally proprietary and expensive with no open-source options now available.

It is strongly felt that any solution that would be put in place to support the JANET community should allow organisations to join without prohibitive entry costs, and without endorsing a particular supplier or manufacturer's systems and protocols.

Further outcomes from the ongoing testing will be available during January 2004.

2. Recommendations

There are a number of key areas that will require further investigation and/or live testing before a wholesale deployment of any solution can take place.

2.1 Authentication and Security

Athens (an Access Management System for controlling access to web-based subscription services) is currently used as an authentication mechanism for the JANET community. This, or any other solution, needs to be integrated with the streaming solutions available to ensure that the system only allows authorised use. This is essential to adhere to content owners' rights licensing conditions.

Many websites and portals to information are protected by Athens logins. However, it is also essential that the back-end systems also authenticate using the same systems to ensure that the servers cannot be fooled into serving content to unauthorised users.

It is recommended that integration of streaming solutions with Athens be investigated further, with the aim of providing a fully authenticated service.

2.2 Cache Solutions

The CDI Trial has investigated Kasenna™ Mediabase XMP SE and Helix™ Universal Proxy which both meet the required standards. However, in the interest of providing a service to a wide and diverse community, it is likely that additional solutions will be required – especially as each has a significant licence fee attached.

Stream caching, in comparison to web-caching, is relatively less well developed. There are fewer solutions in total and there are potentially few, if any, suitable free (software) solutions available (c.f. Squid as an HTTP cache). These options need to be explored and interoperability between platforms assessed.

The current solutions explored have little facility for Real Time Streaming Protocol (RTSP) stream filtering, by URL or by any other means. For a wider deployment, and particularly where FE Colleges and schools may be involved, it will be essential to provide a means of filtering access to content in the same way as access can be filtered using web-caching platforms. The absence of such a facility would be a hindrance to uptake.

Certain platforms do provide such filtering, and these should be explored for deployment where required.

2.3 Digital Rights Management

The issue of Rights Management is more critical in systems where caches may be installed. With the current systems under test from Kasenna™ and RealNetworks®, if content is removed from the origin server a cached version of that file will not be available for viewing from any organisation-based cache. A conversation takes place between the cache and the server when content is requested to ensure that the content is still available and may be served to this user. It is highly likely that certain content providers will require this level of control in any system on which they licence their content for use.

The inclusion of other, third party, cache solutions may break this model of operation and so needs to be investigated technologically, in parallel with a study of likely requirements from rights holders.

2.4 System Monitoring

For any JANET service it will be essential for the service managers to be able to access up-to-date information about the state of the service, and for such things as automatic e-mail,phone or Short Message Service (SMS) messages to be generated on a failure or exception within the service.

Kasenna™ Mediabase XMP provides Simple Network Management Protocol (SNMP) access to server statistics, however, as a standard product RealNetworks®' Helix™ Universal Server does not. Further add-ins are available, but are in a sense 'third-party' products as they come as plug-ins from another section within RealNetworks®. RealNetworks® have indicated that they may be able to give access to these products for testing.

Due to problems within the Kasenna™ product line, Mediabase XMP only became available for testing in September 2003 and the new version is currently being deployed.

It is recommended that service monitoring and control tools are reviewed and tested where possible.

2.5 Pre-loading Content

For lower bandwidth organisations, there is a strong demand to be able to pre-load content onto local caches for local delivery. This is especially critical at organisations connected at 2Mbit/s.

Over a loaded 2Mbit/s link it is difficult to successfully watch, and therefore cache, content during the working day. A mechanism is therefore needed that will allow scheduled download of content.

Whilst this is achievable using the scheduling available on certain players, e.g. from the command line interface to RealNetworks®' client player, these are not wholly satisfactory as the player does not know if a stream breaks up

whilst it is playing. A better solution would be the ability to schedule File Transfer Protocol (FTP) or other reliable downloads from the core servers into the cache. This is possible using the Mediabase XMP solution, but it does require a Mediabase XMP server at both the sending and the receiving point (cache) – which may be beyond the budget of smaller organisations.

3. Background

The delivery of moving image and sound content via Internet Protocol (IP) (streaming) is becoming increasingly common. Significant developments in this emerging technology enables beneficial services to be developed for the JANET community. During the latter part of 2000, UKERNA started planning a streaming content demonstrator project. This planning work identified the Kasenna™ MediaBase technology as the most likely to be able to support the multiplicity of formats and bit rates currently available. As a result UKERNA purchased and deployed a Kasenna™ MediaBase server to gain experience of the technology and its potential impact on JANET.

The JANET CDI requirements study, which concluded during October 2001, confirmed the Kasenna™ MediaBase software running on SGI® hardware technology as one of the leading options that could support a CDI for JANET. Specifically the report recommended that 'A pilot project to serve video content with the Kasenna™ MediaBase product should be carried out.'

In line with the outcomes of the ESYS report UKERNA commenced planning a JANET CDI trial. The objectives of the trial was:

- to assess whether it is feasible and practical to provide a JANET CDI service and to investigate the service delivery and support implications;
- to monitor usage and performance of the different streaming media formats/coder/Decoders (CODECs) and provide guidance and best practice for sites wishing to engage in streaming media;
- to trial Kasenna™ MediaBase's functionality and performance in order to assess its suitability for wider trials and possible national service and to identify alternative solutions for further trials;
- to establish whether the caching of Moving Picture and Sound content is required at the 'edges' of JANET in the context of a JANET CDI Service;
- to trial a variety of hardware in order to establish the best return on investment in terms of price and performance;
- to ensure interoperability with emerging best practice for the creation of other Information Environment projects and initiatives (Managed Learning Environments (MLEs), Virtual Learning Environments (VLEs) and Portals etc).

The trial was structured to involve five participating sites chosen from the HE and FE sectors, together with a lead site that would:

- co-ordinate the participating sites;
- manage the installation and configuration of trial equipment at participating sites;
- gather and collate subjective feedback from sites;
- gather and collate objective monitoring data from the trial's network monitoring system;
- collaborate with UKERNA on the production of a final report.

In order to provide adequate performance monitoring for the CDI trial, a parallel monitoring infrastructure was installed, that provided adequate measurement granularity.

A Call for Participation was published by UKERNA in November 2002 with the trial scheduled to start in January 2003.

Following the publication of the Call for Participation, five sites plus a lead organisation were chosen to participate in the trial. They were:

- Brighton University;
- The Sixth Form College Farnborough;
- Sussex Downs College;
- The University of East Anglia;
- The University of Manchester;
- The University of Wales Swansea (lead site).

The trial was overseen by the CDI Trial Working Group (CDI-WG), that included staff from:

- UKERNA;
- BUFVC/MAAS;
- The University of Wales Swansea (lead site).

The CDI-WG was responsible for providing advice to UKERNA on the direction of the trial The CDI-WG met fortnightly, initially, and then as required throughout the trial by videoconference.

The trial members became part of the Project Group (CDI-PG), that consisted of staff from all of the participating organisations as well as the members of the CDI-WG.

The CDI-PG held face-to-face meetings at the start and end of the trial, and videoconferences were held fortnightly to allow feedback and discussion of issues raised during the trial.

A project manager at UKERNA managed the trial overall, and responsibility for certain elements was devolved, with agreement, to the lead site.

4. Content Delivery Architecture

4.1 Design Factors

It is critical that infrastructure designed for the JANET community considers the education environment. Key differences between the education and corporate network environment, that many large-scale software solutions are designed for, should not be treated lightly.

Whilst this may seem obvious, it is also true that today the largest Content Delivery Networks are those used in the corporate world and those used by cable operators for programme delivery. As a result, the market is geared towards these environments and this can cause suppliers problems when designing, implementing and discussing systems for use in education.

The main issue is generally one of control and management. In most corporate systems, one organisation generally has control and access to all elements that affect the delivery of the service – from the core servers, routers and network links down as far as the desktop PC or, in a cable delivery environment, more often a set-top box.

In an academic environment, the NREN operators, such as UKERNA, manage their networks frequently through Network Operations Centres (NOCs) that may or may not be part of the NREN itself. However, whilst the core routers and links are usually managed in this way, either at the regional or the organisational level, control passes out of the NREN/NOC domain into a more local management domain – such as that of the Regional Network Operators (RNOs). There are also limits to the influence of an RNO over its member organisations, or sometimes more critically, there can be limits over the services the RNO will or can offer even if the organisations are demanding them.

There are severe limits as to how much influence the NREN or NOC can have over the role of a regional or organisational network apart from through co-operation and tightly controlled or monitored Service Level Agreements (SLAs).

Also, in a corporate streaming environment – especially in the delivery of streamed content via cable networks to users – it is often a purpose-built network, designed specifically to deliver voice and video data to a fixed number of subscribers.

JANET, on the other hand, is a multi-service, multi-domain and multi-user environment with many different political and technical drivers pushing various parts of the network in different directions. Whilst applications such as videoconferencing and Voice over IP (VoIP) are increasingly found in corporate networks, applications that require Terabyte data transfers and remote control of scientific experiments seem rarely to be found on networks outside of academia.

In designing a CDI that will support the required numbers of users without losing the ability to deliver high quality content, certain things must be considered:

- the multi-service nature of the academic network with realtime and bulk transfer data in the same pipe;
- the multiple management domains, where access to equipment is not a given from one domain to another;
- that different Regional Networks and organisations will be at different stages of development, and may offer different levels of bandwidth and services such as QoS and multicast to their user communities.

It could be the case that an organisation will want to join a JANET service but the intervening RNO, or may not have the manpower or technological ability to be supportive. In these cases organisations must be able to proceed without RNO assistance.

The infrastructure solution that UKERNA has chosen to trial will allow RNOs, organisations and even individual departments to opt-in to the service with a minimum requirement for assistance from any external organisation.

It should be noted that the infrastructure concept chosen is separate and distinct from any choice of specific hardware or software. It is also the direction that other European NRENs are beginning to explore, though currently they focus more on the streamed delivery of live events rather than video-on-demand as in the UKERNA trial.

The system designed (see Figure 1) consists of a core server or server farm, hosted for optimal network performance at a SuperJANET4 (SJ4) backbone C-PoP. Organisations, or other entities, install a streaming cache at their sites – just as many sites currently run a HTTP or ‘Web’ cache.

Clients that want to view streamed video connect to the core server ‘through’ their own streaming cache. If the video is already cached on the local streaming cache then that content will be served from the local cache, if not, the cache will stream the video from the core server, simultaneously delivering it to the requesting client and storing it to satisfy future requests locally.

Figure 1

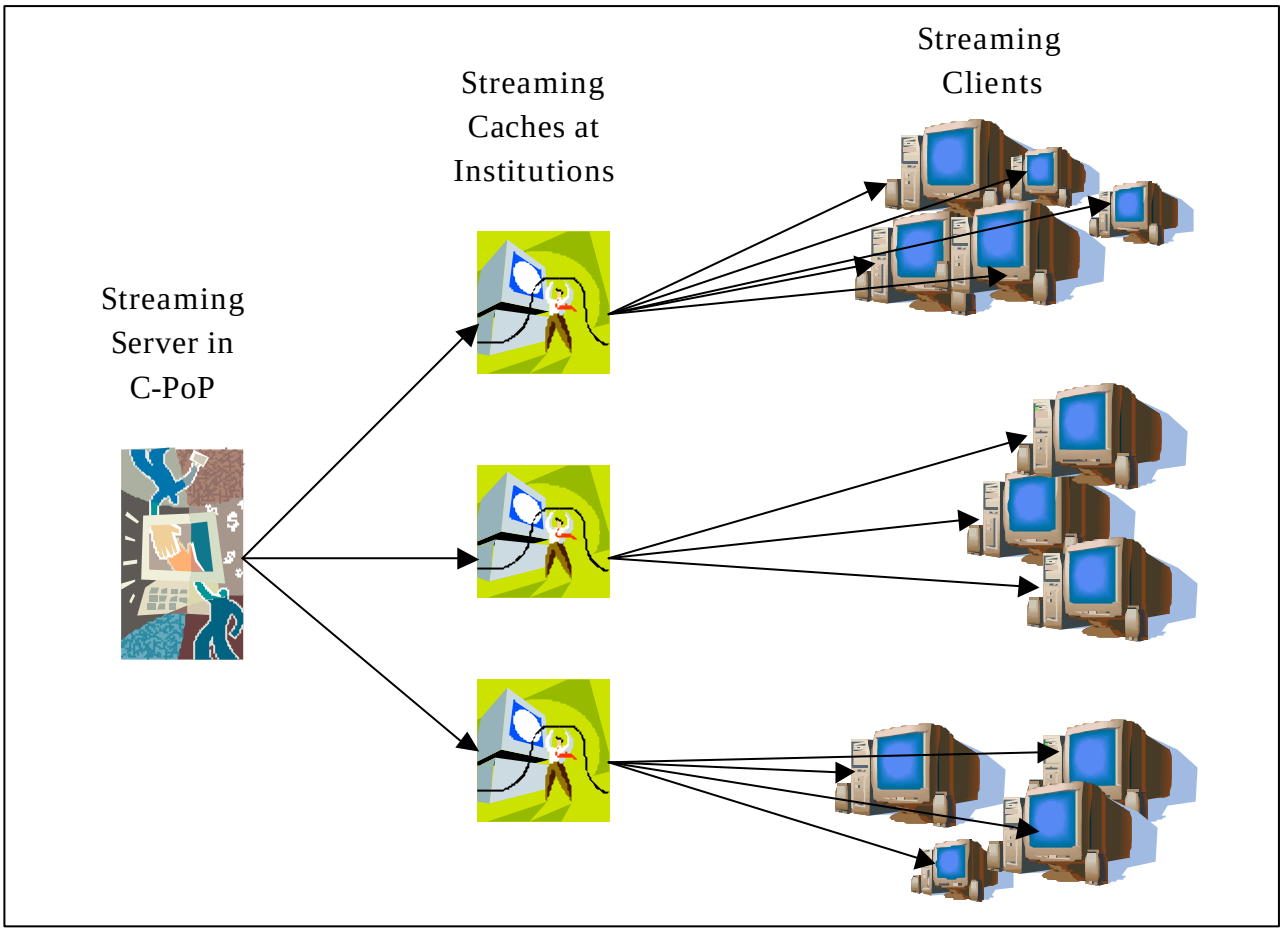

There are a number of benefits to the Server-Cache approach:

- Many sites block unknown traffic at their firewall or Site Access Router. Streaming video like videoconferencing uses dynamically allocated ports, so it can be difficult to wholly secure networks. For all clients in the organisation all the possible network ports need to be open. With a proxy in place all streaming video traffic will come from the outside world only to the streaming proxy and then onwards to the client hosts. This scenario is relatively easy to secure and manage.
- If many clients watch the same streamed video without a streaming proxy in place, the bandwidth used will simply be the equal to the bandwidth of the clip multiplied by the number of users. This obviously scales badly, especially when frequently, as libraries have found with book stock, students tend to want the same things all at the same time. With a streaming proxy in place, once the video has been streamed into the organisation, all other requests are served from the local proxy – minimising the external bandwidth utilized.
- Current low bandwidth sites, such as many of the FE colleges and schools, may only have enough bandwidth to deliver a maximum of a single high quality stream at a time. In this case the streaming cache is not a luxury but a necessity. The ability to then preload content onto the streaming cache in advance of viewings also becomes desirable. If pre-loading is able to be scheduled, it then becomes possible to deliver

high bandwidth streams across very low bandwidth links – providing allowance is made for the transfer time of what may be a multi-gigabyte file. These downloads can be scheduled for networks' quiet time – such as overnight.

The distinction between a proxy and a cache is often mis-represented. For clarity, in this document:

- A proxy is a network device that 'hides' the requesting client's IP address from the server from which it is requesting the data. A proxy has no means of storing data for later local delivery without re-requesting the data from the external source. It operates in what could be called a 'pass-through' mode;
- A cache on the other hand does the same as described above, but also saves to local disk a copy of the data it receives from the external source. It is then capable of passing subsequent requests for data in order to ascertain whether it holds a local copy. If it does it will then serve that data from local disk to the client requests.

Unfortunately the terms proxy and cache are frequently used interchangeably, but this can be extremely misleading. An organisation that purchases a proxy to overcome bandwidth problems may find itself in trouble if that proxy is not also a cache.

4.2 System Description

As recommended by the ESYS report, the trial set out to test an infrastructure based on Kasenna™ Mediabase, which allows an infrastructure as described above to be built. The software also has a large number of additional features built in, such as a simple web-portal to allow users to query the database of assets held on the server, and management and monitoring capabilities. It also allows the scheduled distribution of content from the core servers to caches across the network.

Kasenna™ Mediabase runs on a number of platforms (and operating systems) with the most frequently used including SGI®/IRIX), Sun® (Solaris™ and Linux®) and Intel® based Servers (Linux®). The choice of hardware is flexible within these manufacturers, however in some cases, the choice of operating system version is not.

In order to test performance and scalability and to provide feedback as to the sizing of video proxies, a number of different hardware platforms with different specifications were installed in participating sites. These were:

- 3 x SGI® Origin 300s with TP900 storage units;
- 2 x Dell® PowerEdge® 2600;
- 2 x Sun® LX50.

One SGI® server was used as the core server and was installed into the Reading C-PoP. SGI® platforms were also installed at Brighton and the University of East Anglia. Dell® servers were installed at the University of

Wales, Swansea and University of Manchester and Sun® servers at Farnborough Sixth Form College and Sussex Downs College.

4.3 Installation and Configuration

Equipment was delivered to the University of Wales, Swansea, during February and March 2003, for staging and pre-configuration before delivery to sites.

The SGI® platforms run IRIX, the SGI® version of UNIX®. The Dell® and Sun® server platforms are Intel® based and run RedHat Linux® v 7.3. Kasenna™, unfortunately, (at the time of configuring) only supported RedHat version 7.3 – even though the current version of RedHat Linux® was v8.0. More importantly, Kasenna™ would not support any patches being applied to the boxes – including security updates. In a production network that is connected to the global Internet by high-speed links, the rate of port and vulnerability scanning is such that this approach cannot be recommended. It would only be a matter of a short time before un-patched vulnerabilities would be exploited.

In order to accommodate this, and in order to continue with the Kasenna™ software, the decision was made to lock down the host machines involved in the trial more tightly than was intended. This would have the effect of preventing participating sites from accessing material from servers outside the trial core server until the security of the systems could be ascertained. This would be equivalent to installing a web cache and only allowing access to .ac.uk based domains – secure but highly unpractical for normal use.

The SGI® core server was installed by SGI® in the Reading C-PoP. After installation and SGI®'s initial configuration the lead site inspected the server. IRIX, by default, comes installed with a number of user accounts pre-defined. Unfortunately the passwords are not predefined and are blank. By default IRIX also does not install any sort of firewall. This is how the SGI engineer left the server at Reading – open and highly insecure.

Both these issues, one with RedHat and Kasenna™ and the other with IRIX and SGI highlight the lack of awareness in the commercial sector to the issues of having hosts directly connected to high speed networks that are an integral part of the global Internet.

The Dell® and Sun® servers that came with Linux® pre-installed were re-installed with clean operating systems to minimise security problems and to allow the lead site to document exactly what environment Mediabase required to run.

Installation and configuration of the operating systems and Kasenna™ Mediabase are relatively straightforward, with additional summary documentation provided by SGI® for their hardware – mainly relating to the need to rename disk volumes on the TP900 storage units.

It is worth noting that the Kasenna™ Mediabase software is actually a large and complex package. The installation includes the streaming software, an Apache web server and an SQL database within the software bundle.

As has been mentioned above, there were concerns as to the security and potential vulnerability of the servers in the trial. In the event of a wider deployment, it will be critical that Kasenna™ Mediabase, or any other software deployed, can support up to date and fully patched versions of operating systems. In a normal deployment, the cache boxes must be able to pull content from a wide range of sources – many of which will be outside the UKERNA domain. This is analogous to web proxies, that are able to pull content from anywhere, in order to serve it to local users.

In order to provide information to the participating sites, a password protected website was set up specifically to support the CDI project. This allowed sites to monitor the progress of the trial and allowed the lead site to provide information relating to the network configuration needed at sites.

5. Network Monitoring Infrastructure

5.1 Overview

One of the elements that significantly impact on the performance of streamed media is the underlying network. In the case of the UKERNA CDI trial the performance of the multiple domains (JANET backbone, Regional Network, organisation) that make up the JANET network will have a significant effect on the quality, and therefore the users' perception of delivered video quality.

In order to allow objective conclusions on the performance of the delivery infrastructure itself, a dedicated monitoring infrastructure has been put in place. The purpose of this is twofold:

1) To provide indications of any network issues that may impact the trial and allow the participating sites to correct these to ensure reliable media delivery to their site and;

2) To allow network performance to be decoupled from server performance during the trial.

Item 1) is, in general, good housekeeping and in many cases should be the sort of monitoring that takes place in the JANET backbone, Regional Networks and organisations on a regular basis. In general data on this sort of testing is not released outside the domain doing the testing. For this trial the monitoring spans multiple JANET domains, originating from the Reading C-PoP where the streams are hosted, transiting the JANET core and the Regional Networks and terminating at the organisation next to the streaming proxy and the streaming client. In a way this could be viewed as setting out to prove that the JANET core and the wider JANET network are performing as more generally expected.

Item 2) will allow the decoupling of the server/proxy performance and network behaviour. If media streams are being degraded, is it due to Central Processing Unit (CPU) loading, disk Insertion Order (IO) performance etc? or network-based packet-loss, latency etc.?

5.2 Monitoring System Description

The monitoring system is based on 'appliance' monitoring points – i.e. hosts dedicated to the monitoring task and performing no other functions. The monitoring infrastructure exists in parallel to the streaming infrastructure in order to prevent issues with one system impacting on the other.

In this case the monitoring points are based on Cisco® router platforms, with a Cisco® 7206VXR/NPE400 at the core and Cisco® 805 routers located at participating sites. Cisco®'s Service Assurance Agent (SAA) runs on each of the monitoring points and handles the construction, transmission and reception of each measurement probe.

Additional Unix® hosts gather data from the monitoring points and store it in an SQL database for later analysis. The scripts are custom written in Practical Extraction and Reporting Language (PERL) for the CDI trial and a mySQL database holds the data.

Figure 2 shows the key components of the system: the core and edge monitoring points, and the additional hosts required for processing and storing the data.

Figure 2 - UKERNA CDI monitoring infrastructure

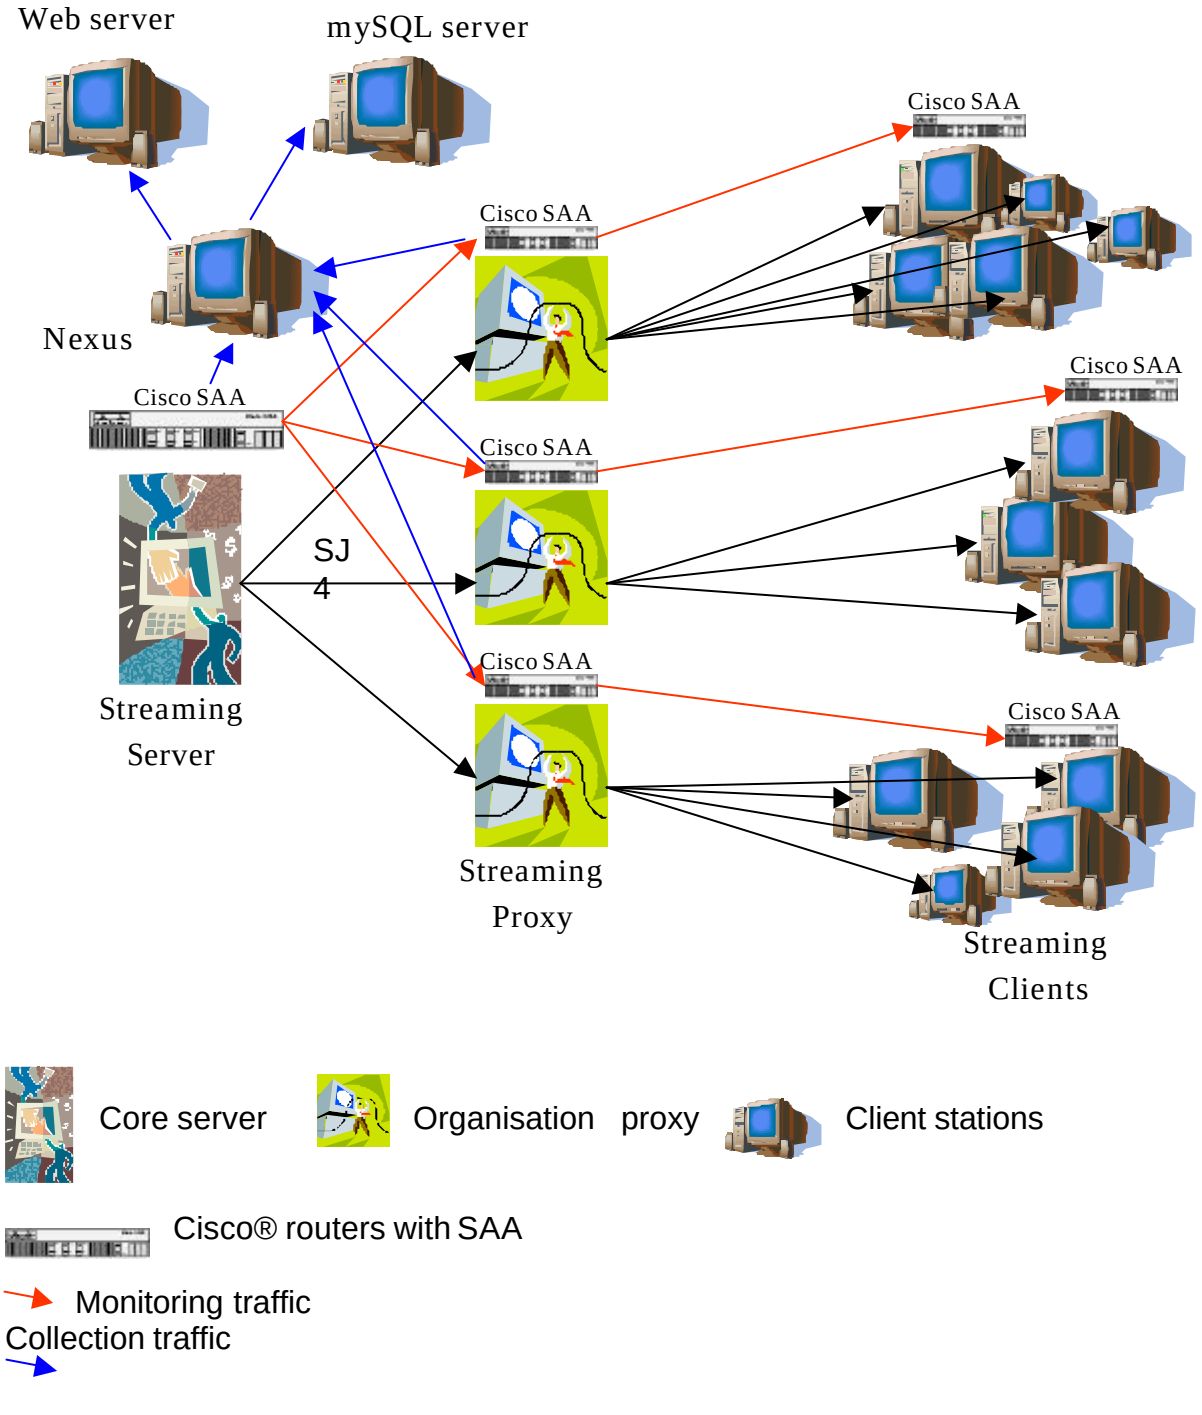

Stream and stream control traffic


Much work has been done, and is continuing, in the European and wider academic and NREN research communities on network monitoring, but as yet there are few precise definitions of what type of monitoring traffic is required to provide statistically reliable evidence of overall network performance. There is a trade off between, at one end of the scale, running a large amount of traffic across a link, which will degrade overall performance for all the link traffic, and running a periodic single probe, that may miss significant outages.

In some cases, when the monitoring is designed to prove reliable delivery of real-time traffic, it has been found that running periodic traffic which bursts a larger number of packets across the network has been a more reliable indicator of network performance than periodic single probes or pings. The model adopted by the CDI trial uses periodic probes of multiple, closely-spaced packets from the Reading C-PoP to each organisation's monitoring points.

Each measurement probe consists of 50 packets sent at 20 millisecond intervals. For each packet data are gathered by SAA on the:

- Round Trip Time (RTT) in milliseconds;
- packet loss (per packet);
- Inter Packet Delay Variation (IPDV) or jitter, in milliseconds;
- One Way Delay (OWD) in milliseconds.

OWD measurements are available if the monitoring router is time synchronized to a Network Time Protocol (NTP) source.

Cisco® SAA gathers the above results from each of the individual packets within a measurement probe and internally summarises them. The output data available for processing from a single measurement probe is:

- the sum of RTT measurement for the 50 packets;
- total packet loss in the probe;
- summary of IPDV for the probe;
- the average OWD.

This summarises the type of information that is available for download from the SAA process following a probe. It should be noted, however, that this amounts to some 42 data per probe. It includes items that describe if the router or SAA process was unable to successfully complete the probe/packet for internal scheduling or other reasons. This enables certain types of spurious data to be filtered out, and can also provide an indication if the router is becoming overloaded with SAA probe handling.

The time that a measurement is made is obviously important to the reconciliation of the objective monitoring statistics with users' feedback forms. For most types of measurements made by SAA, only the time on the source monitoring point is critical, as this provides the timestamp for the measurement packet. This applies to all measurements except for OWD. These measurements require highly accurate synchronisation in order to

achieve reliable results. RTT measurements require the packet to be time stamped only at the source monitoring point and then, having been sent to the edge and the edge returning the packet to the core again, the timestamp and current time are compared to give the RTT in milliseconds. OWD measurements, however, require the packet to be time stamped both at the source, and at the destination monitoring points, then on its return to the source, timestamps are available for both legs of the trip and the OWD for each direction can be calculated. Note that in both cases the resultant data is only extracted from the source monitoring point.

The source monitoring point is synchronised using NTP to provide adequate accuracy for the trial.

The accuracy of NTP clock synchronization has been studied in depth within the TERENA TF-NGN (TF-NGN is a task force composed of representatives of the NRENs and research organisations with the aim of studying and developing technologies that are viewed as strategically important for the NRENs and the GÉANT project as set out in the GÉANT Technical Annex. (year 2 update) and the results have been impressive.) Providing the NTP daemon is configured appropriately and the NTP time-synchronisation traffic is routed symmetrically across the intervening network, accuracy within a millisecond is achievable. If the NTP time-synchronisation traffic passes over asymmetric routes, however, the accuracy can be significantly degraded – the extent being dependent on the relative difference in delays across each of the routes.

6. Content Delivery Trials

6.1 Delivery Platform and Content Format

The CDI trial began implementing an infrastructure based on the findings of the ESYS report. The report recommended the use of Kasenna™ MediaBase Network Edition running on SGI® hardware. Unfortunately, by the time the CDI Trial commenced the Network Edition of MediaBase had been withdrawn, with no product development expected. The product known as MediaBase XMP SE was substituted after discussion amongst the Working Group in consultation with Kasenna™. Mediabase XMP SE provides the same basic features as Network Edition, but lacks support for non-standards based formats. The only file formats that were supported were MPEG1, MPEG2, and MPEG4.

A number of discussions took place on the preferred content format for the delivery infrastructure. It was strongly felt that the format should be open and accessible with no penalties for participation by sites either in cost terms, or by forcing sites to implement a vendor specific solution. In these terms the MPEG standards based formats would be the best option.

The Working Group concluded after research that MPEG1 would be hard to implement due to hardware restrictions, and is seen by many as outdated and coming towards the end of its life. MPEG2, whilst being widely used, was felt to have too high a bandwidth utilisation to be scaleable to a service on the scale of the JANET community. In addition, it was felt that the more efficient MPEG4 Coderr/Decoder (CODEC) would take the place of MPEG2 in due course.

It is worth noting that where bandwidth is not an issue MPEG2 is the most widely used standard for content delivery. MPEG 2 is used in cable networks and in digital satellite and digital terrestrial broadcasting with bit rates of 7Mbit/s and higher.

MPEG4 has been optimised for network delivery, providing the highest quality video and audio at the lowest possible bandwidth and so was a natural choice for the delivery format. Kasenna™ MediaBase XMP SE supports MPEG4.

6.2 Kasenna™ Mediabase XMP SE

Initial testing began on Kasenna™ MediaBase XMP SE while hardware was delivered to the participating sites and installation of software took place. One of the critical aims of the trial was to ensure that participating sites with a low bandwidth connection to JANET could cache content pulled down from the core server at the Reading C-POP on their own CDI Trial installed local streaming cache. This was a fundamental requirement of the streaming infrastructure because low speed (2Mbit/s) connected sites would not be able to take part in the trial unless they could view content. As content was encoded at 768kbit/s, there was little chance that low bandwidth sites would

be able to watch more than one or two streams from the origin server without experiencing noticeable problems with the quality of the streams.

Kasenna™ MediaBase XMP SE required a client-based application, the Kasenna™ Broadband Player (KBPlayer) to be installed on client PCs receiving streams from a MediaBase server.

KBPlayer comes configured to play MPEG1 and MPEG2 streams. In order to play the MPEG4 streams used by the trial it was necessary to further configure it. This meant registering a different Multipurpose Internet Mail Extension (MIME) type, for MPEG4, and effectively telling KBPlayer that MPEG4 content would be viewed using Apple QuickTime® Player. This was used as the KBPlayer was unable to render MPEG4 content. Unfortunately it was not possible to only use QuickTime® Player without KBPlayer due to the MIME types and file formats involved. Documentation on configuring the KBPlayer was poor and required numerous e-mails to Kasenna™ support staff to configure for MPEG4 content.

Using KBPlayer meant there was also an issue in providing access to MPEG4 streams on the Macintosh® platform. Kasenna™ Support informed the lead site that there was no client player application available for any other platform besides the PC. Ideally, any JANET streaming infrastructure would be platform independent, but this was not the case with the KBPlayer.

During configuration and testing the lead site found that the MediaBase XMP SE software did not appear to be able to cache MPEG4 content. This issue was raised with Kasenna™ Support. The absence of any high level Europe-based technical support caused delays during the testing phase. The only high-level technical support available was based in the USA and it was very difficult to get a speedy response to queries, especially when questions or answers needed clarification. It took on average two days to get a reply from Kasenna™ Support.

Kasenna™ confirmed, after a delay of some weeks, that there was a serious flaw in the software that prevented MPEG-4 content from being cached. The lack of a caching function for MPEG4 content made Mediabase XMP SE unsuitable for the trial and the decision was taken by the Working Group to look for an alternative solution that could fulfil the requirements of the trial.

6.3 Helix™ Universal Server

The Working Group surveyed the options for replacement software that would allow the trial to continue as planned. Various suppliers and organisations were contacted and the activities of other NRENs, especially in Europe, were considered.

The Working Group reached the decision that RealNetworks®' Helix™ product line would meet the requirements of the CDI Trial and RealNetworks® UK were contacted and agreed to supply licences and technical support for the duration of the trial. It was noted that RealNetworks® Helix™ product line

could also serve most of the major streaming formats, including MPEG4 and the QuickTime® format.

Due to technical difficulties in obtaining the MPEG 4 encoded material for the trial in a timely manner, and independently to the Kasenna™ Mediabase XMP SE issues, it was decided to proceed with QuickTime® format files which were immediately available for the trial. As MPEG4 content became available it would be added to the available content. Content was available in QuickTime® format as this is one of the formats, along with Windows® Media Player, available from the Education Media On-line service.

Helix™ Version 9 software will not run on SGI® platforms (unlike earlier versions), so new streaming cache hardware was installed at the two sites that had SGI® hardware and a replacement Intel®-based server was installed in the Reading C-PoP as the core server.

The Helix™ software was tested and was found to distribute and cache MPEG4/QuickTime® files correctly. One issue was found within the server software, resulting in files only being cached if they were mounted in a certain place in the Helix™'s file system. This was not a serious issue and was overcome by other means.

The Helix™ Universal Internet Server and Helix™ Universal Gateway software was in general extremely simple to install and configure.

As mentioned above, the only significant issue encountered with Helix™ Universal Server was connected to the caching of content. The QuickTime® Player software used in the trial to play MPEG4 content was configured to use the local CDI Trial streaming proxy in a similar way to that which Microsoft® Internet Explorer or Netscape Navigator® is configured to use a local web cache.

When first implemented at test sites Helix™ Universal Gateway appeared to be unable to cache content. When first tested (pre site deployment) at the lead site, this feature had functioned correctly but upon being deployed on site it did not appear to work. Every time a user selected to play a clip the content was played directly from the core server through the local proxy to the client player - a procedure called 'pass-through'.

Those organisations with a low bandwidth (2 Mbit/s) connection to the Internet were immediately aware of the issue. In most cases they were unable to play even a single instance of a video clip without the quality of the stream degrading to an unusable degree. As has been mentioned previously, this was the expected effect of attempting to stream directly across the loaded 2Mbit/s links, but was confirmed at this stage of the trial.

Although the caching problem was identified immediately, RealNetworks® were unable to identify any solution. Further testing by the lead site and then confirmation by RealNetworks® led to the discovery that content would only be cached if it were distributed from a subdirectory of the main root mount

point on the Helix™ Universal Server. If the mount point was moved using the administrative interface, Helix™ Universal Gateway was unable to cache the content locally. This was overcome by changing the setting of the main root mount point and using symbolic links to point to the file systems where content was stored.

RealNetworks® reported that the bug would be fixed in a subsequent release.

Once the file system on the core server had been re-configured those sites with low bandwidth connections were immediately able to start trials. Those sites with a high-speed connection to the Internet, such as Manchester, were unaware that their content was not being cached, as the access link bandwidth is so high.

Due to the caching issue those sites with slow connections to the core took longer to establish their trials, whereas sites with high-speed connections were able to begin almost as soon as equipment had been installed on site. However, once the low speed connected sites were able to usefully watch content they were able to provide a lot of valuable data even though they were 'playing catch up' with the sites with faster connections.

It should be noted that Helix™ Universal Server and Kasenna™ Mediabase XMP SE function with an entirely different model of distribution. The difference is easiest to explore by viewing what happens from the client player end. With Kasenna™ Mediabase XMP SE, the client requests data from the local streaming cache, and it is the streaming cache that knows if content is available and where that content is located. The client never knows which server the material comes from originally.

Helix™ Universal Server works on the web model, where client players request data from the origin server, but the client player software is configured with the details of the local streaming cache and so will request content through that cache.

The significance of this is that in the Kasenna™ model, the local cache has to be told about all the content available in the system, unlike the Helix™ solution.

As clients request data from the cache in one instance and the server in the other these models are mutually exclusive, and it remains to be seen whether a system based on the Kasenna™' model can access other, external material outside the knowledge of the local streaming cache. This will be tested during the fourth quarter of 2003.

It is worth noting that opinion differs as to which type of distribution model is better. Certainly the Kasenna™ model is highly suitable for corporate content providers in closed networks and cable operators. Whether it is so suitable for open systems that require the ability to pull content from any source, anywhere is unclear as yet. The 'Web' model is more tried and tested in the education world and may prove to be the most flexible in a large deployment.

An early problem encountered at one site during the trial concerned the apparent loss of audio when playing QuickTime® movie streams, and the loss of video with audio only present when rendering MPEG4 streams. When using the QuickTime® browser plug-in the video element of a QuickTime® movie would play back without any problems, however, the audio did not play. Only if the user opened a new copy of the browser window, with the original one still open did the audio then appear from the new window. Similarly, with the MPEG4 content, at first only audio played, however when a second browser window was opened to play the same content using the QuickTime® browser plug-in the video appeared.

This issue was passed back to RealNetworks® Support who were unable to identify an immediate solution. After investigation, it became apparent that the problem lay with certain network ports that are routinely blocked for security reasons at the lead site's firewall. This was subsequently verified by RealNetworks® support. Helix™ Universal Server uses a pair of ports for audio and video delivery and the default port for QuickTime® movie audio was one of the ports that are blocked at University of Wales, Swansea – it is also a well-known port that is used by an audio file sharing system. When a duplicate window was opened on the same PC the port number for the audio part of the stream was incremented past the range of ports that are blocked by the firewall. This was also the case for the lack of video when rendering MPEG4 streams.

Other issues that arose during the trial were due in part to the particular idiosyncrasies of network setups, and the demands placed on the JANET network by the education sector being different to the supposed architectures that the streaming server software was designed to facilitate. In general, streaming software vendors are designing their software to work in corporate Local Area Network (LAN) settings that are by their nature generally closed to the outside world.

In such a setting there is usually only one type of network model in place, one main software platform and one main system administrator taking appropriate precautions to avoid security threats from external sources. At education organisations, especially universities, such control is not always possible or even desirable. Academic network users have very different requirements from their counterparts in the corporate sector. An academic network needs to be as open as possible, will invariably contain a wide range of differing hardware and software platforms and have many possible routes where attacks can be mounted on the network.

One example of how the needs of the academic community differ to that of the corporate sector is the absence of any means of filtering RTSP and Universal Data Protocol (UDP) streams on available streaming proxy servers – a practice widely implemented on web proxies to deny access to websites using the HTTP protocol.

7. Subjective Feedback

In order to provide an effective mechanism for measuring the success or failure of the CDI trial a Subjective Feedback Form was designed with the assistance of staff at the various test sites involved in the trial. Its purpose was to provide an opportunity for the end users of the streaming infrastructure, the staff and students at the test sites, to voice their opinions. Aside from those regarding time and location, the questions that were asked were as follows:

- Title of film;
- Collection Name;
- Did the film run smoothly?
- Were there problems with the sound quality during the film?
- Were there problems with the video quality during the film?
- Did the player lose its connection to the server requiring you to manually restart the film at any time?
- Did you have any other applications, i.e. e-mail, Microsoft® Word, running while you were watching the film?
- Were you comfortable while watching the film?
- Overall, did you feel that the experience was a good one?
- Was there anything else that would have made the experience more satisfying?

The questions were designed to be non-technical and focused on the actual experience of viewing rather than being mindful of any issues regarding speed of connection, client software, speed of client processor, etc.

Overall there were 161 feedback forms completed. The actual number of streams watched during initial testing meant that the total figure for the number of streams watched is actually far higher than 161. However, this figure represents those streams that were watched under the prescribed trial conditions.

The total number of forms completed by each organisation is as follows:

Organisation	**Number of completed forms**
Brighton University	33
Farnborough SFC	37
Manchester University	24
Sussex Downs College	28
University of Wales, Swansea	10
University of East Anglia	29

7.1 Questionnaire Responses

The answers to the questions can be summarised as follows.

Title of film:

The two most watched media were the titles ‘An Introduction to NMR Spectroscopy’, from the Sheffield University Learning Media Unit with 34 viewings and ’79 Springs’, from the Educational Television Films collection with 16 viewings.

Collection Name:

The three most watched collections were those from Educational and Television Films, Films of Scotland, and the Sheffield University Learning Media Unit respectively.

Did the film run smoothly?

125 respondents answered ‘yes’, and 36 answered ‘no’. Interestingly, there was no correlation between network size and available bandwidth as to whether the answer was yes or no.

Were there problems with the sound quality during the film?

Possible answers to this question were: Never, Frequently, and All the time. Of the 161 forms returned 112 respondents did not experience any problems with audio whilst watching.

Were there problems with the video quality during the film?

103 respondents from all test sites stated that there were never any problems with the quality of the film.

15 respondents commented that they frequently had problems.

5 reported that they experienced problems all the time.

Did the player lose its connection to the server requiring you to manually restart the film at any time?

143 respondents reported ‘no’ to this question
18 reported ‘yes’

Did you have any other applications, i.e. email, Microsoft® Word, running while you were watching the film?

16 responded ‘yes’

145 responded ‘no’

The most common applications that were running at the same time were Microsoft® Internet Explorer, the application used to choose the content in the first instance, and either some form of antivirus agent running in the background or an e-mail client such as Microsoft® Outlook.

Were you comfortable while watching the film?

It was interesting to note the answers to this question as the repurposing of computer labs as places to view and ingest learning material is a significant issue for organisations wishing to utilise streaming technologies for learning. Simple things like providing headsets to listen to streamed content, ensuring sound cards are connected correctly and earphone sockets can be accessed without having to remove any anti-theft devices are all serious issues for organisations and support staff.

Of the respondents who replied 159 answered ‘yes’ to whether they were comfortable. Only two answered ‘no’ giving as reasons the busy ambient noise in the computer lab and poor quality headphones.

Of those who responded ‘yes’ the majority commented on how well lit and quiet the environment was. A few also commented on the high quality of the displays used, a testimony to the increasing quality of computing facilities both in the FE and HE sector.

One factor to do with accessibility does stand out from the comments made. This is concerned with how simple it is for the user to adjust sound levels on their PC or even to discover if a sound card is present. Upon reflection it is clear that organisations will need to address the issue of educating users in managing audio on the PC, perhaps through the use of guides or introductory sessions.

Overall, did you feel that the experience was a good one?

129 users responded ‘yes’, 12 responded with a ‘no’ answer. In most cases the ‘no’ answer followed either a break up of the stream or they did not have enough time to finish watching the whole film. Users that answered ‘yes’ often stated that what they had watched to be highly informative and an excellent learning resource – a reflection on the high quality of the material available.

Was there anything else that would have made the experience more satisfying?

32 users responded ‘yes’, 129 responded ‘no’. Those replying ‘yes’ were often merely commenting on the fact that the stream broke up during viewing. Other suggestions for improvements were as follows:

- disabling of screensaver – disrupted the viewing experience;

- better choice of films;
- better quality of audio on films – lots of hissing;
- would have liked video window to have been larger;
- metadata made available for media;
- better quality of films – users found them uninteresting or badly conceived;
- interactive activities to complete while watching the film, such as having to answer questions posed by the film.

It is worth noting that the specification of client machines used to play trial content has not been an issue - considering the high load rendering video can put on processors. Most, if not all, of the machines used at participating sites are more than capable of playing many instances of streamed video. During visits to trial sites to install server equipment the lead site team was in general impressed by the high specifications of the client PCs available for use as part of the trial.

8. Monitoring Results

The Sixth Form College Farnborough has a 2Mbit/s access link to JANET. As the content available through the CDI Trial was encoded at 768kbit/s, it was suspected that Farnborough would be dependent on the streaming cache-to-cache incoming media for local playback. Otherwise, at best they would be restricted to being able to play back a single stream on a single computer at any one time.

Figure 3 shows monitoring data from a 'normal' college day. No streaming is taking place at this time.

Figure 3 - Farnborough data 6 May 2003

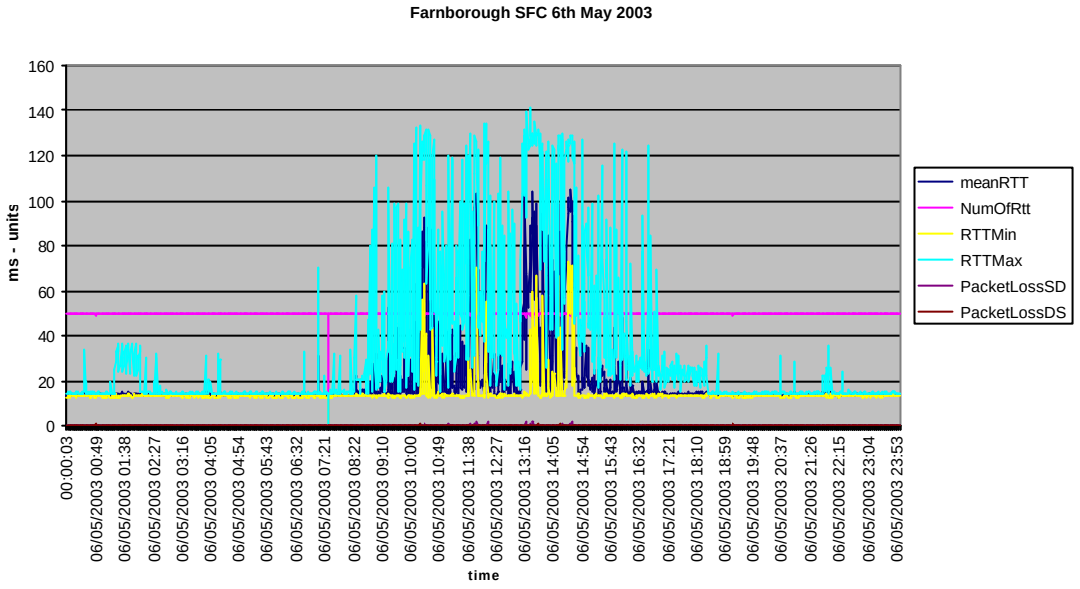

The graph shows that, while the network is fairly reliable at packet delivery (in pink), the maximum delivery time (in light blue) for those packets varies substantially during the day. Significantly it follows the working day very closely, with an increase in delay after 08:30 but dropping again after 17:00.

The minimum delivery time (in yellow) increases rapidly at certain points in the day, indicating a significant load on that link. The pattern of latency shown along with the slight increase in packet loss, shown by the small deviations of the very bottom lines on the chart (in black and brown) seem to indicate that the network link is struggling to keep up with demand at peak times. And while packet buffers in the routers are ensuring delivery of most packets, the fact that they are also being overloaded, causing tail-drop at times, suggests that Farnborough would likely see problems with streaming traffic during daytime. This has been borne out in practice by reports from Farnborough as a result of watching content.

If we compare Farnborough to Manchester for the same period, in Figure 4, we see that the pattern is quite different. This was expected as Manchester

has a high-speed connection. Generally the difference between the minimum and maximum RTT is only a few milliseconds throughout the day. The variations seen early in the day, between about 07:30 and 08:30, show that use was almost certainly being made of the JANET 'at-risk' period.

Figure 4 Manchester University data 6 May 2003

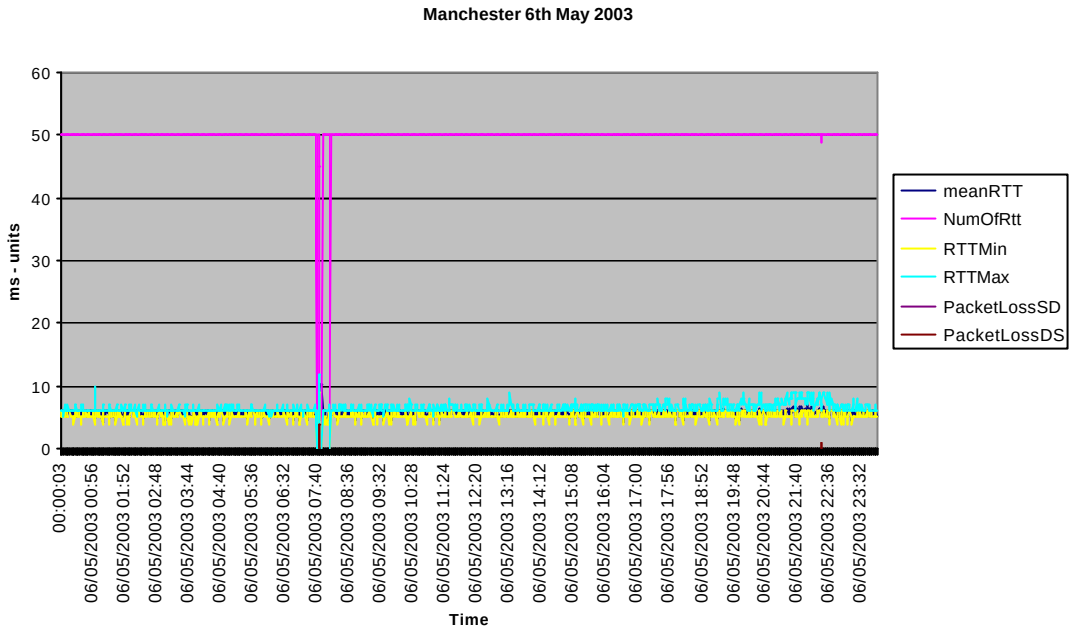

This is the profile of a healthy network, with only a single lost packet during the 24-hour period (ignoring the at-risk period).

If we look at another example of the same day, in this case Swansea, Figure 5, we see that while there is no significant variation in packet delivery times, there are a number of packets lost in transit across the network, in both directions.

Figure 5 - Swansea data 6 May 2003

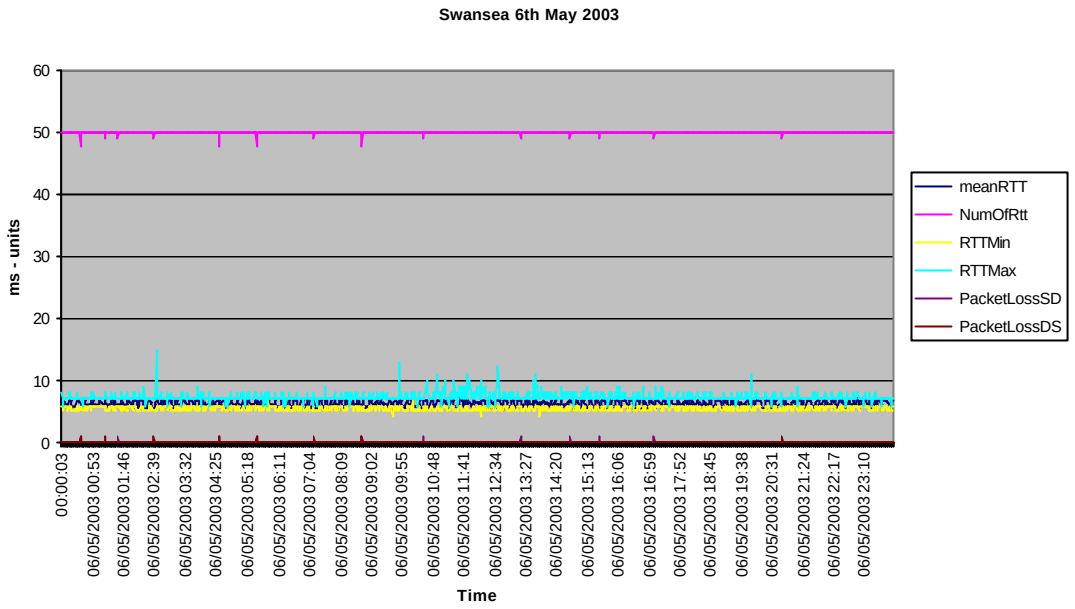

Whilst the lost packets account for only 0.03% of packets in the sample, it may indicate an area worth exploring, as there are few reasons that traffic should be lost. It is the case however that Swansea have the streaming cache and its monitoring point in a temporary location while their machine room is 're-jigged'. Once this has been done the equipment will be directly connected to their core campus switch.

With the above loss level Swansea have not reported any problems running multiple (10-30) streams into the campus.

All the above data are representative of the networks during a 'usual' day. For the purpose of this trial we will now move on to look at data from days when the organisations have been running CDI Trial tests – and see if there is any significant difference during streaming media.

8.1 Initial Testing

On 8 May both Manchester and Swansea ran their first content delivery tests. Between 10:00 and 11:00 Manchester ran approximately 10 simultaneous streams at 700kbit/s each. At this time all streams were being delivered directly to the client machines from the core server – no caching of streams was involved.

The graph shows that there was no significant effect on the network to Manchester University at the time of the trials. The disruption at 15:03 is unrelated to streaming trials.

Figure 6 – Manchester data 6 May 2003

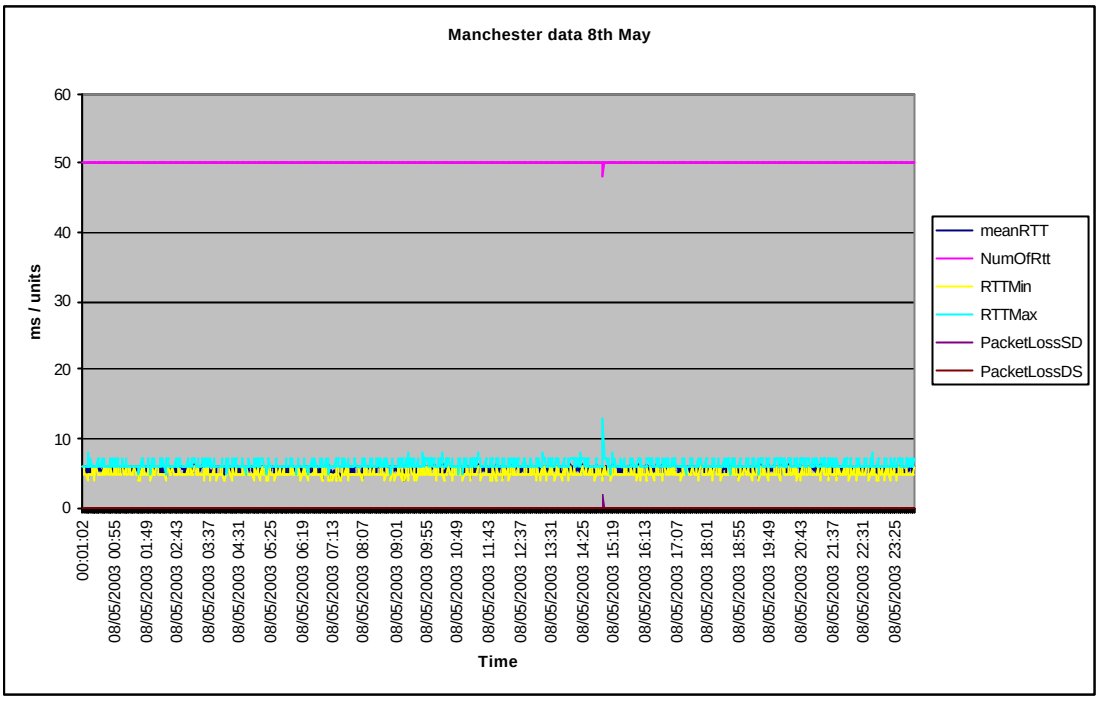

Data for Swansea also shows no significant deviation from a normal pattern at 15:00 whilst tests were being run. This test also involved 10 simultaneous streams.

Figure 7 Swansea Data 8 May 2003

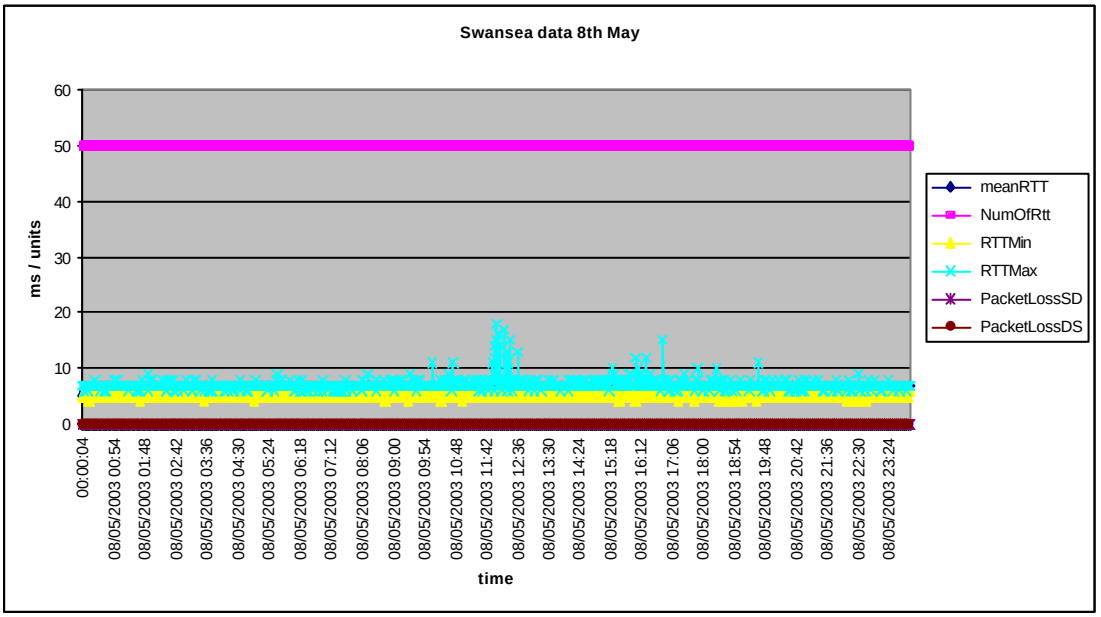

8.2 Conclusion

From the samples above, and the balance of the data from the monitoring system, both the JANET backbone core and the Regional Networks involved are performing well with minimal disruption to packet delivery.

There have been some occurrences where the monitoring system has shown abnormal packet loss or delay, and these have mainly been down to the organisation's networks. An example of this is the University of East Anglia (UEA), whose packet delivery latency for a period of time was noted to exceed that of sites connected at far lower bandwidths. Figure 8 shows the pattern of delay over a five -day-period.

Figure 8: The University of East Anglia (UEA) – High Latency

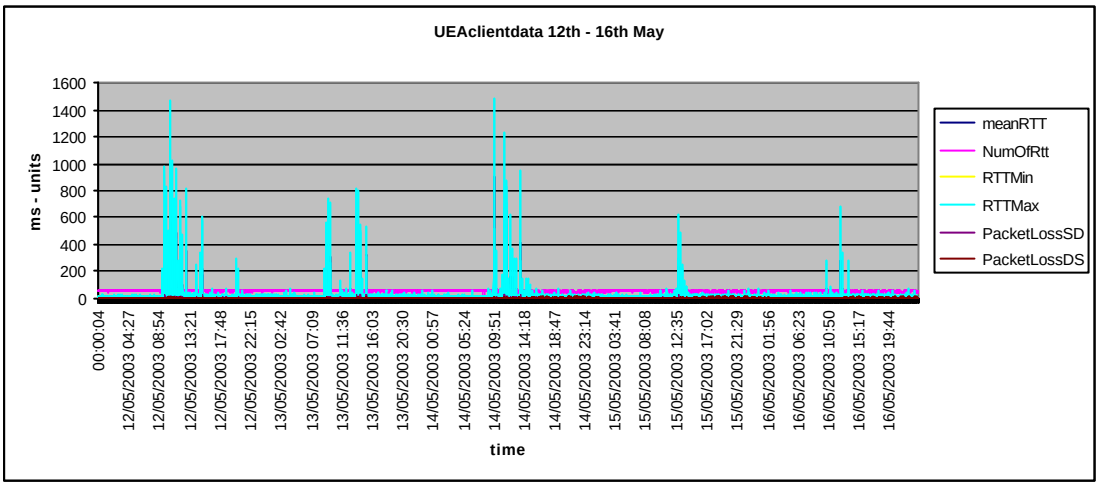

It is significant that maximum packet latency in this graph is shown to be peaking at 1.4 seconds, relatively long in network terms. It is more usual in JANET connected networks to see packet loss than delays of this duration due to the limited length of packet buffers in the sort of routers associated with high speed networks. In this case, a firewall installation was suspected, though being specified to have a high throughput may have been buffering significant amounts of data, leading to the latency figures seen.

It is also worth noting that subjective feedback from UEA indicated that content delivery did not appear to be unduly suffering. This was probably due to the relatively large buffering, (as much as five to ten seconds) done by streaming clients to allow for minor network interruptions.

9. Overall Conclusion

It is widely acknowledged that the potential for the use of streamed media in education is tremendous. There are only two significant elements in achieving uptake, that of content quality and in provision of suitable access to that content.

The quality of material that is available, from those sources used in the trial and also from sites themselves – whether that be 'home-grown' material or recorded off-air under agreements (such as the one with the Education Recording Agency), is excellent.

This CDI Trial has shown that it is possible to deliver high quality content effectively over the education networks in the UK and with a degree of reliability that was not available before – due to network upgrades and higher bandwidth provision than has previously been available.

It is however the case that most involved in implementing large scale architectures of this kind, i.e. involving cache infrastructures, do not feel that either the marketplace, or the technology has yet fully developed; software is continually being updated, fixed and having features added. This is the best period in which to be exploring such areas, as it is now that the JANET community has the ability to influence the direction and structure of future products, and in general, vendors appear to be interested in what the community has to offer.

The recommendations proposed in this report aim to explore the areas that are key to adoption by the JANET community, dealing with issues such as access and authentication, security, interoperability and rights management.

The ongoing phase of trials will complete in December 2003, at which point a further update report will be published. Whilst that report will not provide the answers to these questions, it is hoped that the direction that research and deployment must take will become clearer.

The most significant issues may well be the interlinked ones of rights management and interoperability. It is certainly the case that a caching system could be deployed today that would serve the needs of the community to view content whilst ensuring the rights of the content owners are protected. However, that system would be proprietary and potentially attract a large licence fee – a significant disincentive. It would not interoperate with other vendors' products or have any 'free' software options for sites.

The ongoing challenge is to find a solution that will enable multiple vendors' equipment to achieve the same features, giving flexibility of choice to sites wishing to participate in streaming.

10. Acknowledgements

UKERNA wishes to thank all those at the participating sites who have been involved in the trial.

Also thanks are due to the BUFVC and MAAS for advice and content encoding.

Trademarks