CASE STUDY

February 2009

East Midlands Broadband Consortium

The UK Access Management Federation
FOR EDUCATION AND RESEARCH

Interviewed: Professor Peter Thewlis, embc Technical Consultant
Ben Ellis, Technical Director, Synetrix

BACKGROUND

embc Procurement Ltd, formerly East Midlands Broadband Consortium, is a local authority company wholly owned by Derbyshire, Leicestershire, Lincolnshire, Northamptonshire, Nottinghamshire, Rutland, and the cities of Leicester and Nottingham. It is one of the 13 regional education networks in the UK, and one of the 10 in England. In terms of geographical spread, if not population, embc covers around one sixth of England, stretching from the North Sea to Derbyshire and from just south of Grimsby almost to Milton Keynes.

When embc began in 2000, says Prof. Peter Thewlis, unlike many RBCs there was little to no broadband infrastructure in the region. The consortium was therefore able to take a different route to many of the other RBCs and build a dedicated network right across the East Midlands.

The embc network now connects directly to 2100 schools in eight local authorities, with 300,000 registered users. Physically, it is a 1Gbit/s core network with two Tier 1 Data Centres that offer all services to all sites except Lincolnshire. The Data Centres are connected on a 10Gbit/s ring. Thus the network is fully resilient, both for connectivity and between the two data centres. If one Data Centre becomes unavailable, the services are automatically delivered from the other.

embc uses a range of different technologies to provide connectivity to schools with bandwidth scaled to meet the demand. Secondary schools are currently being moved to 100Mbit/s connections. If the average traffic on any link goes over 60-70% of capacity it will normally get upgraded, following agreement with the site.

The network is an open MPLS[1] network running across the whole of the East Midlands. It operates as a single carrier class infrastructure over which networks and VPNs are delivered, including the common schools curriculum network or the corporate networks of different authorities. Northamptonshire, for example, has five or six different networks (VPNs) running over the embc infrastructure. Thus the infrastructure can be used to interconnect with county council wide area networks, enabling access to other sites including children's centres, adult learning centres, libraries, museums and so forth.

The Data Centres offer a full range of school specific, safe, secure and appropriate ISP services to education, including filtered internet access, e-mail, website hosting and videoconferencing. Access is through a SharePoint 2007 portal with instances for embc and each of the local authorities. Peter Thewlis explains that whilst the portals include My Sites and other services for each authenticated user, it is not a virtual learning environment. These are provided either locally or through the member LAs, meaning that numerous different VLEs are in use across the region, some with single sign-on through the appropriate portal and others with integrated services.

Default Internet filtering profiles can be set up for each site in the region and individual profiles for each user once they have authenticated. This enables a restrictive level of filtering for non-authenticated users with less restrictive filtering levels being given to authenticated users as they progress through school.

1 Multi-Protocol Label Switching

The open network concept means that whilst embc supplies the infrastructure and core services, it also enables third party services and applications to be delivered to connected edge sites; for example, remote network support for school networks, CCTV, alarm monitoring and building control. A supplier to one site uses the network to configure and operate access control to the buildings remotely. ACLs[2] in school routers control access across the network, ensuring that one school can't see servers or equipment in another unless embc configures it to do so. This enables the embc network to be configured to support multi-site working and support across school clusters and partnerships.

embc manages an IP addressing structure for the regional network. Private IP address ranges are allocated to the various networks across the region and for each of the local area networks that connects to the infrastructure. Only if a school wants an Internet-visible server does a public IP address need to be allocated from the embc ranges. This means that all devices on the embc network operate on a private IP address range. A user can log in at any point on the network and be given their allocated user profile and filtering level.

This is where Shibboleth starts to become important, as it can perform non-IP based authentication not only down to a site level but to individual users without the site or user having a public IP address.

A REGION-WIDE IDP

All the East Midlands local authorities agreed the need for a regional IdP service. embc is registered as an IdP[3] with the UK Access Management Federation and is now the IdP for the East Midlands region. At the moment this is only for schools but it will expand beyond that. Ben Ellis of Synetrix was responsible for its implementation, using the Active Directory (AD) created for the embc services. Some local authorities and school clusters want their own ADs and work next year will enable these to be federated with embc's central AD.

Everything hinges on identity, says Ben. A user just has to login once to get e-mail, a portal site, a personal web site and access to services provided by Shibboleth-enabled Content Providers. Having one central AD means that embc can ensure the right attributes are available through the IdP service. Through the embc SharePoint portal, schools can create logins for all their pupils, allocating them services and setting their filtering levels.

Ben observes that interesting issues arise with VLEs on the network, as they tend to need more data than is possible under standard eduPerson attributes. This has been recognised in the Shibboleth community for a while and will be resolved on a case by case basis. However, if a school or local authority contracts a third party to provide a VLE, there are data protection issues about the release of information which must be handled by the school or local authority themselves. This places the responsibility at the correct level; for example getting parental agreement prior to the release of pupil data is the responsibility of the school or LA and not embc, which is not the appropriate organisation to handle these issues on behalf of schools.

In collecting the original data for the IdP, Ben says Synetrix has observed a few issues. For example, the recent changes to the synthetic eduPersonScopedAffiliation highlighted the issue that some education sites, usually education centres of some sort, have had pseudo-DFES codes issued that only exist in the local authority. The new synthetic scope extension has rules for this kind of establishment but many of the sites have historically had numeric codes rather than the alphanumeric codes now required.

Getting the data right is the biggest challenge in setting up an IdP service, Ben says: do that right and nearly everything else falls into place. The 'out of the box' Active Directory schema had to be extended to hold the required eduPerson attributes and a further extension of the LDAP[4] schema to hold ukSchoolperson attributes, a proprietary extension which was developed before the UK federation decided on metadata

2 Access Control Lists
3 Identity provider
4 Lightweight Directory Access Protocol

standards and is used for access to internal embc services that have been Shibboleth-enabled. This extension holds data that cannot be held within eduPerson schema: for example, whether the user is a school network administrator.

User data in embc is managed using Microsoft ILM[5] and a proprietary web-based user management system. As a new identity is set up, e-mail accounts are provisioned and e-mail and Internet filtering policies are set. Any errors in the data processing are highlighted for manual intervention and correction by the school administrator. Ben says this system works very well even with the 300,000 individually registered users – a number rising by a further 30-40,000 a month as more and more schools take the service up.

Peter Thewlis points out that it is in content and service providers' best interests for them directly to control and manage who has access to their content. They can directly determine such matters as free trials, subscription levels and discounts with their customers. embc does not need to be involved in the process unless a different attribute release policy is required for the service. Either the embc or the Local Authority can add links to the service or content on the appropriate portal site.

OBSTACLES OVERCOME

Synetrix originally intended to install the Shibboleth 1.3 IdP but instead went directly with Shibboleth 2.0. At the time this was not supported by the UK federation and so the metadata used 1.3 endpoints and the IdP works with both Shibboleth 1.3 or 2.0 service providers.

An ISA[6] array acts as the Shibboleth WebISO[7] as well as the authentication mechanism and SSO[8] for the Microsoft technologies, Sharepoint and Exchange e-mail. The Shibboleth IdP was originally designed to sit behind the ISA array; however, it was found that the Shibboleth 2.0 IdP could lose the session information due to the SSL setup this would require. To solve this, Synetrix configured the Shibboleth infrastructure to sit outside the ISA array. Ben expects this set-up will change with future IdP upgrades.

Ben also touched on the difficulty of load-balancing the IdP service. Currently a active/passive setup is in place which is also expected to change with future upgrades.

Drawing from the embc project and other Synetrix implementations, Ben noted that the Shibboleth middleware has not reached the stage where the technical knowledge of the people implementing it is at general sysadmin level. To implement Shibboleth, implementation staff need to understand more than basic Java, Tomcat, Apache, programming environments, Secure Sockets Layers and SAML. This, he feels, holds back implementation of Shibboleth generally.

The IDP servers in embc are deployed in a virtual server environment. One of the issues in implementation was with the clock synchronisation – the virtual machines tended to race ahead or slow down in time, which is problematic in a Shibboleth environment. Apparently it took a software fix to the Virtual Machines software to resolve.

SHIBBOLETH TAKE-UP BY USERS AND PROVIDERS

Access to JISC Collections has been piloted in the region and embc is talking to other content providers but Peter Thewlis observes that it is still early days for the use of Shibboleth in the schools sector. To date few content providers in the school sector are using Shibboleth and similarly there are few organisations (RBCs or LAs) providing comprehensive IdP services. However this is starting to change rapidly. Now that the regional IdP service has been set up, Shibboleth-enabled content or services can be sold to every school in embc and they will have single sign-on to that content or service.

5 Identity Lifecycle Manager
6 Internet Security and Acceleration
7 Web Initial Sign-on
8 Single sign-on

What is badly needed, they both say, is Shibboleth exemplar content for primary and secondary schools that is free to use. Whether this content is a good sample or a cut down version of a service is irrelevant as long as it is something that says 'This is what Shibboleth can do for you.' This will really assist Shibboleth in taking off in the UK.

Our thanks to Peter Thewlis and Ben Ellis for agreeing to be interviewed for this case study.