WHITEPAPER: Web Filtering for Schools – The Requirement for Differentiated Filtering.

AUTHOR: John Hackett for E2BN.

The Background

There are two recent reports which have reflected upon the use of the internet in schools and, in part, make reference to filtering issues. The (2009) Ofsted Report "The safe use of new technologies" and the (2008) Byron Review.

The Ofsted report makes a distinction between "locked down" and "managed" systems (see paras 6 and 7). With a locked down system "... almost every site has to be unbarred before a pupil can use it..." while managed systems "... also have inaccessible sites, [but] there are fewer of them". The exact nature of a "managed" system is not explicitly defined in the report but our reading of this distinction could be characterised as the difference between "block unless requested and vetted" and a presumption to allow a site unless there is clear reason for it to be blocked.

The Ofsted report makes very clear that the managed approach is to be preferred with the explicit recommendation that schools should "manage the transition from locked down systems to more managed systems to help pupils understand how to manage risk; to provide them with richer learning experiences; and to bridge the gap between systems at school and the more open systems outside school" (page 6).

As stated in the Byron Review:

> "There is a general social consensus, reflected in our approach to film and television content, that explicit pornography and violent material such as videos of executions is not suitable for children. However, there is no such consensus about material such as non-pornographic nudity, violence or death in an educational context (such as information about wars or the holocaust) and the websites of extremist political parties. Similarly, many parents would wish to *stop young children from stumbling across such material*, but would be *keen for their children to see such material when they are older teenagers* or *when it can be put in an appropriate context*." (4.56)(our italics)

and

> "The decision about what constitutes "inappropriate content" can be highly subjective. What one person views as harmful, another might find offensive, whilst yet another might see it as an important, empowering learning experience for their child; *and this view is likely to change depending on the age of the child*. An example of this might be a sex education website. In consequence, any attempt to block content which falls into these grey areas would leave some parents unhappy that the system was either too restrictive or not restrictive enough ..." (4.57) (our italics)

While the Byron Review is mainly concerned with internet access and filtering in the home these same issues – of inadvertent access to inappropriate material and the important of age differentiation in what is considered to be inappropriate – occur even more clearly in the school context where pupils from 5-18 as well as staff are using the internet and any filtering system used must be able to cater for these diverse needs.

Particular Needs of Staff users

Teaching staff have very particular needs when considering a web filtering system. For example, if staff have the same filtering rules as their students then they are unable to use their professional judgement to evaluate blocked sites. They may feel that a particular site is suitable for their sixth form students but not for others: but how are they to make this judgement if the site is blocked not only to all students but staff as well? This is not to say that staff should have completely unfiltered internet access but there should be a presumption of allowing staff access to sites which fall within the "grey areas" noted by Byron.

Staff need to research their subject and may want to use certain sites which are, for whatever reason, blocked to students.This is independent of whether they may subsequently want their students to visit the site. To use the example cited in the Byron Review (4.56) above a History teacher may want access to a site about the holocaust to gather information and images for a classroom lesson. The teacher may, indeed, feel that the site is not suitable for student

viewing but they still need access themselves in order to contextualize the information it contains.

Similarly there are clearly situations where it makes sense for staff to be able to show to a whole class parts of a particular site that they would not want pupils to be able to use unsupervised. For example, there are many very useful videos on YouTube that staff may want to show to a class. But there is also a vast quantity of very inappropriate material there so it is generally blocked to students.

Requirements

Taking these two reports into consideration allows us to be able to specify in broad outline the facilities required for a school web filtering system:

(1) The general "ethos" of the system should be to "allow" a site unless there is some reason to block it. The reason may be a generic one (i.e. block all online games sites) or very specific (e.g. this site contains language which is not suitable for Year 9 students). The key is that there is a defined policy of when a site is, and is not, allowed.
(2) It must be possible to have age differentiated filtering rules for students. As we see from Byron (4.56 above) a particular site (or even a particular page within a site) may be considered suitable for pupils in the sixth form but not those in Year 9 or younger. An example could be a site designed for lawyers or law students where criminal cases are discussed in great details with crime scene photographs. This would clearly be distressing if viewed by a primary school pupil but could provide valuable research to a sixth former studying A-level psychology or law.
(3) There must be a set of filtering rules designed specifically for staff use.

The first of the items above is about the way the lists are managed while the second two are concerned with how these lists are applied to a particular user. In particular it makes clear the need for the filtering system to be able to apply different filtering rules (rulesets) to different populations of users within each school. As an absolute minimum each school needs to be able to access two contrasting rulesets: one for its students and one for its staff. For a more flexible approach and to allow for age differentiation within the school a wider range of rulesets would be preferable. For example, a secondary school could use one ruleset for the lower school, another for GCSE students, another for sixth formers and one for the staff.

Some schools also like to have the ability to restrict internet access further as a disciplinary measure: rulesets that deny all internet access or allow only very restricted access (a "walled garden") are seen as very useful in this context.

Architectural issues

Given this need for multiple rulesets some architectural issues need to be addressed.

The first, and perhaps most fundamental, is that filtering must take place **before** caching. This is required because both the user and client workstation identity are unavailable to the filtering service when the web request has been passed through a cache and so only a single ruleset can be applied. Also, due to the caching mechanism some 25-35% of all requests will never be examined by the upstream filtering system as they will be met from the local cache. That is to say the result of a single web request to the filtering system will often be sent to multiple users. Due to these factors the possibility of any age discrimination is lost if any caching is implemented before the filtering.

So, if caching at the school is required (e.g. to increase performance on constrained bandwidth), then the filtering needs to be implemented in the school, either on the cache hardware or via a separate filtering service accessed prior to the cache.

It is also possible to have this age differentiated filtering when the filtering is provided as a central, local authority wide service. All that is required is some way of directing a user's browser to the filtering system in such a way that the ruleset to be used can be determined. There are a variety of ways in which this can be achieved: (1) using multiple ports on a common proxy IP with each port relating to a specific ruleset, (2) using multiple IP addresses for the service where each IP provides a specific ruleset, (3) using a LA wide authentication system to assign users to rulesets at login. In the first two cases the school's own network system would be used (via group policies in AD for example) to provide the appropriate proxy details to each user's browser.

All the methods above provide a common collection of rulesets and URL lists for the sites using the filtering system. It is also possible to provide a more tailored service where the schools themselves have some control over the filter rulesets and lists. In architectural terms this can be provided either as a central service or via local servers.

A central service could be provided as, for example, an “farm” of virtual filter servers hosted on a central array of physical servers with all the advantages of central management. Such an architecture would allow each school to have direct access to the management of the filtering service they receive (via a web interface) in order to make specific changes to the rulesets and lists without affecting other schools in the authority. The disadvantage of this remotely provided local service is that caching cannot be done locally: each browser request must be handled by the remote virtual server before any caching can take place.

Using a local server (either real or virtual) has the added advantage of allowing caching to be done within the school which reduces the bandwidth requirements – typically in E2BN we find that about 25-35% of requests are “hits” on local caches.

There are several critical advantages to using a local service (whether provided locally or remotely)

(1) By linking to the local directory services the filtering server can be configured to use an appropriate ruleset for each user, add usernames into the filter logs and give school staff access to reports of user activity.

(2) The url lists can be adjusted to more closely reflect the circumstances and ethos of a particular school rather than relying on a “one size fits all” set of rules and lists managed centrally.

(3) Any changes to the lists can be made instantly and without reference to a third party who may, or may not, want to make the change for all the connected schools.

E2BN's Protex Web Filtering System.

The Protex Filtering System was designed in conjunction with schools in the E2BN region to address just these issues. It provides a wide range of age differentiated rulesets (profiles) with the ability to be used both as a central array of high availability servers or as a local filtering service. Local systems are available as either a dedicated hardware solution (filtering and caching) or as a Virtual Machine (filter only) . All types of system can be integrated into a user authentication system such as MS Active Directory or Shibboleth. In this case particular filter profiles can be assigned to groups or user attributes within the authentication system.

Both the local and central systems allow for the local modification of the E2BN managed lists so that the filtering provided by the Local Authority or at school level can be tailored to their users.

Links

E2BN:	http://www.e2bn.org/
Protex website:	http://protex.e2bn.org/
Protex system types:	http://protex.e2bn.org/cms/products.html
Protex Documentation:	http://protex.e2bn.org/protexdoc/toc/index.html
Byron Review:	https://www.education.gov.uk/publications/standard/publicationDetail/Page1/DCSF-00334-2008
Ofsted Report:	http://www.ofsted.gov.uk/resources/safe-use-of-new-technologies