Web Filtering for Schools

Introduction

The Internet is, without doubt, one of the great achievements of the 20th century. Created in the 1960s, the Internet initially only connected a few university computers. Later, in the 1980s, Tim Berners-Lee developed the World Wide Web while working at CERN to allow researchers to collaborate more easily. It was later given to the world with no restrictions and, since then, there has been an explosion in interconnected computers ("the Internet") enabling users to communicate with each other. To share. To learn. To teach. To collaborate.

In 2019 there were over 4.5 billion internet users - nearly 60% of the world's population. (See: https://internetworldstats.com/stats.htm).

There are, however, consequences of this ability for billions of people to interact. Not all internet users share the same political, moral, religious, or sexual outlook. Nor do they share the same opinions on censorship. Some think anything, however outrageous or extreme, should be available to everyone and that all filtering is censorship, which they reject. Others go to the other extreme and would censor anything which threatened their particular world view. As usual, most people's opinion lies somewhere in the middle: exactly where depends on the users of the system being filtered.

That very extreme views and images are available, and easily found, on the Internet is undeniable. That these sometimes appear in the results of legitimate searches is also true. The value of a good filtering system is that it can protect users from coming across objectionable material, freeing them to use this amazing tool with confidence. As we enter the twenty-twenties, 40 year's on from the Web's invention, and 20-years since Google's inception, it is easy to become blasé about the wonders of the Web and too easily scared of its darker side.

Many organisations from small companies, to ISPs, and even whole countries use web filtering systems: this document will concentrate on what to look for in a web filtering system designed for school use.

What is Web Filtering

What is a Web-filter?

Before we consider what Web Filtering is, we need to define some terms. In particular, we need to understand that “the Internet” is the collection of **Inter**connected **Net**works. When you connect your home computer to your ISP (Internet Service Provider) you are connecting to their network which connects to the larger Internet. By itself, the Internet just allows computers to connect to each other - it provides no useful services on its own.

A wide range of services run over the Internet - the most common ones being email and the World Wide Web (WWW) or just the “Web”.

In order to manage the traffic of each of these individual services the connection between any two computers is divided into ports. You can think of a port as a lane on a motorway with each service sticking to its own lane. There are standard ports (the so-called “well-know” ports) for all the common services - there is a list here - although these are not always used.

So, the Web is just one of the many services that run on top of the Internet. Its standard ports are 80 (http) and 443 (https) but web-sites may also use other ports for a variety of reasons. We will assume the standard ports unless otherwise stated. For most filtering systems the actual port being used is not important, as long as the traffic is directed to the filtering system and conforms to the standard Web protocols.

In the context of web-filtering ports play an important role. For example, you can prevent all web-traffic leaving your network (by blocking ports 80 and 443 at the firewall) except when it comes from the web-filtering server.

The basic idea of a web filtering system is to prevent “inappropriate” pages - the definition of which will depend on the context - from being viewed by users of the protected network. Web filtering software only works in response to requests made by users to read specific web-pages: it is not a firewall, nor is it anti-virus software. Some products many include these facilities but, in this guidance, we will disregard them as they are not generally considered part of the web-filtering solution.

Similarly, web-filtering is specific to Web (e.g. http/https) traffic and will not affect other internet services - e.g. video-conferencing, chat, messaging services, etc. It should also be noted that many mobile Apps are not web-based and connect to the internet over ports other than 80 and 443. Some Apps do, however, use these ports to connect to the Internet (as they are often allowed out through local firewalls) but use non-standard protocols. This can cause problems as the web-filter may intercept the traffic, but be unable to respond correctly to the App.

Email is a somewhat special case in that a web filter will not filter standard email (POP3, SMTP) but will affect web-based email (e.g. Googlemail, hotmail, yahoo mail) as these services use the standard web protocols and will be filtered in the usual way.

Why filter the web?

Having clarified what a web filter is, it is time to consider why one is required and, in particular, why having a robust web-filter within schools is an essential aspect of safeguarding both staff and pupils.

Even within the education community there is a range of opinions regarding the use of web-filters. How strict should they be? What should, and should not, be blocked? There are some generally uncontentious categories (porn, gambling, extreme violence, for example) but others will be more nuanced, and different schools will have different opinions: for example, drugs, extreme right/left wing political views, games, etc. This is where disagreements are likely to arise which will need sensitive handling by the administrators of the filter.

There are some opinion formers who see web filtering in schools as bad idea and would prefer to see all school networks unfiltered. This argument usually takes a form similar to this:

"Swimming pools can be dangerous for children. To protect them, one can install locks, put up fences, and deploy pool alarms. All these measures are helpful, but by far the most important things one can do for one's children is to teach them to swim." (From National Research Council study, "Youth, Pornography, and the Internet*)*

Superficially this seems to be a good argument, but it is setting up a false dichotomy. Of course it is important to teach swimming, but that does not negate the need for the other precautions.

What it also fails to recognise, to continue the swimming analogy, is that you do not teach children to swim in a stormy sea but in a controlled, calm, pool where they can be taught in a safe environment, with lifeguards on watch. The same applies to navigating the web: teach about internet safety in a controlled environment where the dangers are, as far as possible, removed. They will then be more prepared to negotiate the internet at home which will, almost certainly, be less well filtered.

Before moving on to discuss the specific requirements of schools, there are some general principles that any web-filtering system should demonstrate.

- Minimise overblocking (i.e. blocking of pages which should be allowed)
- Minimise underblocking (i.e. allowing pages which should be blocked)
- Provide a means for users to alert the system administrators of an over/underblock and receive a speedy response.
- Regular updating of URL lists

School duties and compliance

Before looking at the particular needs of schools, and the constraints they put on a web-filtering system suitable for them, we will consider their safeguarding and "Prevent" duties and the role filtering plays in carrying them out.

The SWGfL have produced a web-page (http://testfiltering.com) to check that the filtering system you have in place is using two key lists which will go some way to ensuring compliance with these duties: the IWF (Internet Watch Foundation) list and the Counter-Terrorism Internet Referral Unit (CTIRU) list.

Safeguarding

The Department for Education's statutory guidance Keeping Children Safe in Education (2019) obliges schools to ensure that "*children are safeguarded from potentially harmful and inappropriate online material*" and, therefore, that governing bodies and proprietors should "*ensure appropriate filters and appropriate monitoring systems are in place*" (Para 87).

However, these filters must be managed so that over-blocking does not lead to "*unreasonable restrictions as to what children can be taught with regard to online teaching and safeguarding*" (para 90).

So, while schools should use filters to safeguard their pupils, the Department do recognise the balancing act that is required to block out inappropriate content while keeping to an absolute minimum, the blocking of sites that are appropriate and useful for teaching and learning.

In fact, they go further in the **Annex C: Online Safety** section of the guidance by recognising that the definition of "appropriate" will involve local judgement and is not an absolute: "*The appropriateness of any filters and monitoring systems are a matter for individual schools and colleges…*" (page 97). The department refer to the UK Safety Internet Centre's guidance on what is appropriate, and that schools have obligations under the Prevent Duty which also need consideration.

Prevent Duty

All schools also have a legal duty under the 2015 Counter Terrorism and Security Act which gives legal force to the "Prevent" strategy published by the Government in 2011. The aim of the Prevent strategy is to reduce the threat to the UK from terrorism by stopping people becoming terrorists or supporting terrorism.

All specified authorities (which includes schools and local authorities) must comply with this duty and will be expected to maintain appropriate records to show compliance with their responsibilities and provide reports when requested.

We note here some key points and the Prevent Revised Duty Guidance gives more detailed guidance.

When carrying out this duty all the specified authorities must have "*due regard to the need to prevent people from being drawn into terrorism*" (para 26) and, of particular relevance here, will include "*considering whether **IT equipment available to the general public should use filtering solutions** that limit access to terrorist and extremist material*." (Para 45, our emphasis).

Sections 57-76 relate specifically to schools. In particular, schools are "*expected to assess the risk of children being drawn into terrorism, including support for extremist ideas*" (para 67) and "*expected to ensure children are safe from terrorist and extremist material when accessing the internet in school,* ***including by establishing appropriate levels of filtering***." (Para 71, our emphasis).

Particular Needs of Schools

Web filtering in schools is an inherently more complex problem than in a business where users are all adults and can be filtered in the same way.

This is not an acceptable approach in schools where the filtering needs to be appropriate for the changing needs of children as they develop and for staff who need to prepare lesson materials. Schools are under a duty to safeguard their users while, at the same time, providing access to the broadest spectrum of materials available online to support teaching and learning. Their web filtering solution must be evaluated with both these, somewhat contradictory, aims in mind.

A school web-filtering system does not reduce the importance of educating all users (both children and staff) in safe and appropriate online behaviour. This will equipping them to use the web on devices, and in locations, where the safeguards provided in school do not apply.

Studies by Ofsted ("The safe use of new technologies" (2010)) and the psychologist and child safety expert Tanya Byron ("Safer Children in a Digital World" (2008) commissioned by the Department for Children, Schools and Families) both focussed on the importance of teaching children to manage risk in this area, empowering them to become confident and safe users of the Internet.

The key idea is the *management of risk*: it cannot be eliminated, but neither should it be ignored. To offer acceptably open access to staff and students, manage the risk of coming across inappropriate material, and provide a safe environment for students, it is essential that a school-based web filtering system should have the facility for multiple filtering levels. The absolute minimum is two: one for staff and one for students. A wider range of age-appropriate levels would be preferable, however, as the needs of primary age pupils are very different from those of GCSE or sixth-form students.

Further issues arise around the range of devices to be filtered and their ownership. School-owned computers (desktops, laptops) can be locked down so that users cannot change the proxy settings to bypass the web-filter. Mobile Device Management (MDM) software should be used to manage all phones and tablets allowed to connect to the network and, to ensure maximum coverage, the MDM should be able to handle multiple platforms.

It is essential that all web-page requests are made via the web-filter. How this is managed technically (see *Enforcing Filtering*, below) does not affect the filtering policy or configuration but it may affect how it is deployed on the local network.

Types of Web Filter

Web-filters can be situated either on the local network, or off-site. Both have advantages and disadvantages.

Before discussing these different locations we will outline the basic filtering process.

When a user requests a web page, it can be directed to the web-filter in a number of ways. The browser may refer to a locally configured proxy or use WPAD to discover the proxy automatically. If transparent filtering has been implemented the browser will just make the request via the default gateway (usually a firewall or router) which will redirect the request to the web filter. This is not an ideal solution - there can be issues both with https traffic and mobile apps using ports 80/443 but not standard web protocols - but it has the advantage that the client does not need any special configuration. This is especially useful for guest devices. The same network can use multiple methods: they are not mutually exclusive.

Once the web-filter has received the request, a number of actions may be taken depending on the filter's facilities (see *Filter Methods* below) but in each case the request is either allowed or denied. If denied, the filter sends a page back to the client indicating that the request has been blocked and why. Otherwise, the filter requests the page and returns it to the user.

Remote Web-filtering Service

With a remote filtering solution, all requests for web pages are sent to a "cloud based" server which does the filtering. Possible advantages of a remote filter are:

- the "heavy lifting" of the filtering process - particularly if content filtering is being used - can be run on hardware scaled to meet the demands of the filtered network,
- only the initial request and acceptable pages are transmitted along the school's internet connection which may save bandwidth: blocked pages are never transferred across the local internet link.

The disadvantages are that **all** requests must be transferred over the school internet connection, even those for URLs which will later be blocked, and there is no opportunity to make use of local caching (see below).

How the advantages and disadvantages balance out will depend on the profile of internet use at a particular site.

It is critical, however, that if a remote system is chosen it still allows for local customisation (e.g. adding or removing URLs from the URL lists) and user based filtering (e.g. having a different set of filter rules for staff and students). Ideally, a remote filtering system should be linked to the local network's authentication system so that changes to the user database are reflected automatically in the filtering.

Local Software only

It is possible to download and install web-filtering software (commercial or open source) and provide an "in-house" solution but this is not recommended for anything other than home use or experimentation. Filters are not trivial to configure and manage: supplied lists will need frequent updating, and commercial lists are unlikely to be tailored for education.

In addition, filtering can be a very CPU intensive operation, so optimising the hardware will have a major impact on its performance.

A better solution is to use a virtual appliance (e.g. a Virtual Machine with filter software pre-installed, configured and optimised). Small schools can run this appliance on a spare local server. Large schools may already have a virtual server platform (e.g. Hyper-V, VMWare, etc.) to manage the appliance.

As noted above, regarding locally installed software, the filtering performance of a virtual appliance will depend upon the host server's specification and available resources. These will need monitoring to ensure acceptable performance.

Local Appliance (Hardware and Software)

While a virtual appliance has the advantage of being less expensive, it is reliant for its performance on the underlying hardware. A hardware/software combination, fully installed, configured, and optimised will provide consistent performance scaled to the school's requirements.

In many cases such an appliance can be supplied as a fully managed service with remote management of software and updates to the software and lists.

Where on the network?

If a local system is chosen then where should it go on the network? There are three main options: as just another local server, as a device "in front of" the firewall/gateway, or as the gateway itself.

If the filter is a hardware server, then it can be used as the default gateway with routing and firewalling capacity. In many ways this is the ideal solution as it makes any attempt to "get around" the filtering harder as the only way out to the internet is via the filter server. It is possible to install a Virtual server and use this as a gateway, but this is likely to be less secure than using dedicated hardware.

It is not always possible or desirable, however, to replace the exiting gateway. Placing the filter on the network so that the only way to the internet is via the filter provides the same obstruction to users trying to bypass it while only minimally disturbing the current network architecture. In this case all non-web traffic will be passed through to the gateway while web traffic will be intercepted by the filtering software.

Finally, it can be treated as just another server with an IP address on the LAN. In this case each client must either be configured to make web requests via this IP or some transparent redirection must be in place at the gateway/firewall to redirect web traffic back to the filter server. To prevent the filter being bypassed, the gateway/firewall should block all web ports (typically 80 and 443) access outbound from all but the filter server's IP address.

To cache or not to cache

One advantage of filtering locally rather than remotely is that a local filter can make use of a local cache to reduce the bandwidth usage of the local internet connection. Pages downloaded by the filter can be cached locally, so any subsequent requests for the same page made via the filter can be fetched from the cache rather than the original server.

How much this will impact on bandwidth usage, however, will depend on the web traffic profile of each site. As more sites move to https (which cannot be cached), the value of a local cache will decline. Our current experience is that some 60-80% of web traffic is now https and this will only increase over time.

Filter Methods

Url Filtering

This is the simplest form of filtering: the URL of the request is matched against lists of URLs, and if there is a match, the appropriate action is taken.

Filter lists are generally broken into categories (for example adult, pornography, gambling, social networks, malware, etc). By combining these categories in various ways users can be provided with a tailored filter experience by, for example, allowing staff to use social networks but blocking them to students.

The default action of a URL filter is to allow the contents of the URL to be returned to the user: only explicitly listed URLs will be blocked.

While terminology may vary depending by product for filters which only do URL filtering, the largest list in each category will be the “blacklist”. A match to one of the URLs listed there will cause the page to be blocked to the user.

There may also be a “whitelist” in each category: a match here will allow the page to be sent back to the user even if it would, otherwise, be blocked by the blacklist. For example, suppose *example.com* is blocked. Adding *example.com/allowthis.html* to the whitelist would allow this one page.

The exact nature of these lists, their syntax, their sizes, whether they are pre-filled, automatically updated, or user modifiable are all product dependant.

Before committing to a particular product the way the lists are supplied and managed, how changes are requested, and how much local control is available via a local management interface all need to be considered.

Finally, it is important that any web-filtering system should deploy the IWF (Internet Watch Foundation) list, the Home Office PREVENT list, and the Police IWL (Infringing Website List).

On-the-fly categorisation

As we have seen above, a traditional URL filter will compare the requested URL against the installed URL lists. If a match is found, the appropriate action is taken. If there is no match, the usual default it to allow it. If a “walled-garden” approach is taken then the default is to block it.

Some systems will, in the case of a no-match, employ an algorithmic approach (often using some flavour of AI or machine-learning) to categorise the page. It is this categorisation which then determines whether the page will be allowed or not.

Categorisation may be a real-time process, or may take place after the fact as a way of dynamically building the URL lists. If a visited site is subsequently categorised as, for example, “malware”, then later visits will be blocked. The effectiveness and accuracy of this automatic, hands-off process will, of course, depend on the machine-learning algorithm.

Web Reputation Filtering

Reputation filtering takes an entirely different approach and may either be the main filtering process or only used to filter any non-matching URLs.

The reputation of a site is, typically, generated algorithmically but may also involve some human interaction. Each is assigned a reputation rating (from, say, 0-100) and the filter rules will block all requests which score over some given value.

As with Categorisation, above, the effectiveness of this approach will depend on the quality of the rating system, the algorithms behind it, the size of the reputation database, and any latency in obtaining the result from the centrally held database.

Content Filtering

In all the methods above, the decision to block or not is based on the URL being requested. This may be by direct matching to an existing list of URLs, or by being assigned a category or reputation based on prior analysis. Importantly, they are not based on the actual content of the page being retrieved.

This can be a problem as dynamic, database driven sites, may supply differing results for the same request. Also, with very large sites (Wikipedia, for example) listing, categorising or assigning a reputation for every page becomes an impossibility.

Content Filtering overcomes these limits by analysing the contents of the actual page retrieved and giving it a score, or weight. This is then used to determine whether it is blocked or returned to the user.

The weighting of the contents can be very processor intensive and is usually used only for those requests which do not match any in the URL lists. The content checking process may be relatively straightforward or highly complex. It will certainly include looking for matching words and phrases, but may also include matching on regular expressions (in both the content and URL) and specific signature strings.

Content filtering of http (i.e. un-encrypted) traffic is straightforward in that the page is retrieved by the filter in plaintext and can then be analysed before being sent back, or not, to the user.

With https, the situation is more complex. In the case of a traditional proxy an encrypted connection is made between the client browser and the target site which is passed through the proxy: the proxy software cannot see either the individual URLs subsequently being requested nor the pages' content, so no filtering is possible.

In order for the content of https sites to be filtered the proxy must be able to insert itself between the client and the target site. This is **https interception** and can be controversial (see here for example). For schools, however, the needs of child protection make the ability to inspect https traffic a fundamental requirement. With suitable configuration any privacy issues can be addressed by, for example, whitelisting banking sites so that they are not intercepted.

The main difficulty with using https interception is that a special SSL certificate needs to be installed on every client: this can be managerially complex, but tools are available to automate the rollout.

The process of https interception is that the client makes an https request, which the filter responds to by setting up an encrypted connection between the client and the filter. This is important as it maintains privacy over the local network. The filter then makes a connection to the target site using its own certificate, again maintaining privacy over the internet connection. The filter can now decrypt the responses and analyse the content is the usual way. If the response is acceptable, it re-encrypts the content and returns it to the user. If not, a page explaining why the page has been blocked is returned instead.

The encryption/decryption processes are very processor intensive, so the platform the filter is running on must be scaled appropriately to handle the expected number of requests.

Enforcing Filtering

A web-filtering system can only be truly effective if all web traffic is directed through it. This can be achieved is several ways which should be used in combination to ensure the best protection for the users.

Network Operating System: For devices owned and controlled by the school, the first point of control is when a user logs in to a device. At that point the network can set the proxy for various services - e.g. http, https, FTP, SOCKS, Streaming (RTSP), etc. The http and https proxies need to point to the web-filter server and users prevented from modifying them later.

Firewall: The network may set the http/https proxy and prevent the user changing it, but what if an application ignores the operating system level proxy and either tries to go direct or use its own proxy? While not common on desktop operating systems, some mobile Apps do. In these cases the way to prevent unfiltered access is to use a local firewall and block any outbound access to ports 80, 443 (and any others identified as accessing the web) from any IP other than the web-filter itself.

WPAD/Pac file: Where it is either not possible or desirable to set a client's web proxy via the network operating system, the WPAD protocol may be used instead. This has the advantage of being able to control which proxy is used in more complex ways. WPAD is an accepted protocol understood by a wide range of platforms and is often the simplest way to set the proxy in a heterogeneous network supporting different device types.

Transparent Proxying: This technique redirects traffic destined for the web to the web-filter without making any changes to the client itself (but see below). The default gateway (a router or firewall) redirects any traffic on port 80 or 443 (and possibly others) to the web-filter rather than forwarding it to the true destination. Transparent filtering works well for http traffic but can cause issues with https unless full https interception (see *Content Filtering* above) is used. However, as noted previously, https interception requires a trusted certificate to be installed on each client which may not be acceptable on guest devices.

Bypassing Filtering

Assuming the network has been set up to enforce filtering in the ways described above, then how can the filtering still be bypassed?

Proxy sites: the most common way for circumventing the filtering is by using web sites running proxy scripts. Typically, these have a text box where you enter the URL you want to visit. The site then acts as a proxy - making the request on your behalf and displaying the result.

These proxy sites can cause major problems for URL only web-filters as they will only be blocked if they are explicitly listed. The sheer quantity of such sites - there are many thousands of them with new ones being created every day - makes it impossible to list them all.

Proxy sites are, arguably, less of a problem for Content Filters as the contents of the page returned may still cause unsuitable requests to be blocked.

Content Filtering software, as well as looking for words and phrases, may also look for signature character strings in the proxy home page itself. There are various popular proxy scripts which can be identified in this way and blocked.

VPN (Virtual Private Network): While proxy sites are web-sites accessed via a browser in the normal way, VPN client software installed on a device sets up an encrypted connection to a remote server. Once set up, the browser on the local device will make requests over this encrypted connection bypassing any local filtering server.

There are many companies offering this service (VPNs do also have legitimate uses) for a few pounds a month.

The VPN client will usually set up the encrypted connection over one of the standard ports (e.g. 80, 443, 53, etc) as these are frequently open on the firewall. To block the unauthorised use of VPNs, the firewall can be configured to block these ports unless they come from specified IPs.

No school-owned device should have VPN software installed. In fact, as a matter of policy, no non-admin user should be able to install any software at all! Guest devices may try to use VPN software so all other means should be employed to prevent VPN software running over the network: by blocking ports at the firewall, blacklisting domains providing VPN services and, where known, the IP addresses of their servers.

Summary

The internet is a superb resource for teaching and learning. Since the development of the World Wide Web in the 1980s the Internet (and, more specifically, the Web) has become a powerful educational tool: from YouTube videos to web-based learning platforms. The undeniable benefits of access to the web do, however, come with a downside. Amongst the educational, the informational, the mundane, the humorous, there is also the anti-social, the fake and the extreme. Web filtering is required to protect the children and young adults in school from exposure to the worst of the Internet..

Web-filters are not only used in schools - businesses use them to restrict users' access to non-business related sites, ISPs offer them to home users, and some government use them to control their populations. The latter, in particular, fuels the debate on whether web-filters are ever acceptable. We would argue that, in schools at least, they are as child-protection and network security are far removed from overt censorship and suppression.

Schools also have special requirements with regard to web-filtering. In particular, the wide age range of users in a school means that "one size fits all" filtering is not acceptable. The needs of staff and students are very different, as are those of primary age pupils and GCSE or sixth-form students. Any web-filtering product designed for schools must cater for these differences.

The safeguarding and "Prevent" obligations of schools require that the IWF and Counter-Terrorism Internet Referral Unit (CTIRU) are deployed by their filtering system. The SWGfL's test page is a useful resource for checking that this is the case.

It is important that the filtering system processes all web page requests: how this is enforced and where the filtering is actually done will vary between products. Web-filtering can be supplied as a remote service (i.e. where requests are sent to a server on the Internet for testing) or by hardware or software installed locally.

Setting up the network to enforce the proxy to cover all possibilities can be a complex task, so other precautions should also be employed to prevent users from circumventing of the web-filter. A firewall can block all outbound access on ports 80 & 443 except from the filter server's IP address, for example, which will also reduce the opportunity for VPNs to be used.

Multiple techniques are used to filter web requests: different products may use any or all of them. The simplest is the straightforward blocking of the actual URL requested by matching it against a list of "bad" URLs. Two other techniques based on the URL involve some level of dynamic filtering using AI or machine learning to either categorise the URL (e.g. as porn, gambling, etc) or give it a reputation score. These are then used to decide whether the request should be allowed or blocked. Finally, there is content checking which does not rely on the URL but the actual content of the page once it has been retrieved. This is the most computationally demanding method but is also the most effective as the actual page that will be seen by the user is being checked rather than just the URL. Each of these techniques can be used in isolation or combination.

Any potential web-filtering system for schools must also have a robust way for local users to make modifications. For some products with a single, global URL list this may be just a means of submitting change requests which may, or may not, be acted upon. In others, local changes are possible so that the school can diverge from the default list as they see fit. These aspects of the product need careful consideration before being selected.

Recommendations

Any school based web-filter should provide:

1. Some element of differentiated filtering. As a minimum, it must offer two levels: one for staff and one for students. Ideally, it should offer more levels to cater for different age groups.

2. The ability to integrate the filtering system with the local network authentication mechanism so that users and groups can be assigned to a particular filter level without undue management overhead.

3. A range of filtering techniques (as detailed above). Our view is that a school web-filtering system must provide content checking of all URLs not explicitly listed.

4. Because of the rapidly advancing use of https it is now essential that the filtering system can handle https interception in order to view the URLs being requested and their contents.

5. Users should be able to modify the filtering system so that their school's particular ethos and requirements can be met. For example, the system should allow schools to whitelist specific URLs that are normally blocked (and *vice versa*).

6. In addition to (5) there should be a method for schools to make the filtering service provider aware of uncategorised, or incorrectly categorised, sites for the benefit of others. For example, by submitting the URL of an unblocked pornographic site for inclusion in the global blacklist.

7. Any system should include fully populated URL lists and pre-configured content filtering word and phrase lists which must be regularly updated (ideally at least daily).

8. Any system must be deploying both the IWF and CTIRU lists to comply with the school's safeguarding and Prevent obligations.

9. While not usually included in the filtering system itself, its deployment should include firewall/router configurations to mitigate the ability to bypass the filtering system (e.g. port blocking and preventing VPN clients from being installed).