Guide to filtering on Janet

Guide to filtering on Janet

Title:	Guide to filtering on Janet
Reference:	*GD/NOTE/006*
Issue:	5
Document owner:	John Chapman, Director information security policy and governance
Authorised by:	Jisc legal
Date:	28 March 2025
Last Reviewed:	21 March 2025

Document control

1. Owner updated. Reference to Online Safety Act added. Specific legislation removed from Section 7.

1 Introduction

This document is intended to provide guidance for organisations connected to the Janet network on the subject of Internet content filtering. It outlines what filtering is performed by Jisc and the approaches that can be taken by individual organisations that wish to apply further constraints on the material accessible to their users. It also points to other sources of advice.

Filtering can be used for two distinct purposes: to protect customers' equipment from external attacks via the Internet and Janet, and as part of customers' response to inappropriate activity

– however the customer may define that – by their own users. Jisc's policy is that the former (security) requirement is likely to be common across all customers and can largely be implemented using existing network equipment: customers would prefer this protection to "just happen", without their needing to act. The latter (policy/content) requirement may well vary both between and within customer organisations: they are likely to want – in some cases be required by law – to help the person engaging in the inappropriate activity, so any filtering technology must fit into their local processes for providing that support. The Janet connectivity service is therefore designed to implement security measures based on the best current information available to Jisc, but to let each customer organisation choose its own content filtering policies and methods. Implementing these controls at different technical levels should avoid conflicts: in particular customers can choose – without interference from the network – whether to implement their content policy using local filters, systems elsewhere on Janet, or elsewhere on the Internet.

2 Background

The Janet network is an integral part of the global internet, connecting together the Local Area Networks (LANs) of education and research organisations across the UK. With user organisations ranging from schools to advanced scientific research groups, the range of requirements of the network is very wide. The fundamental purpose of Janet is, therefore, to provide all connected organisations with highly reliable, high-speed, national and international connectivity that enables customer organisations (individually or in groups) to build or obtain the particular combination of services they need. Unlike consumer Internet Service Providers (ISPs) Janet does not routinely provide higher-level services, such as domain name resolution (DNS) or content filtering, to all customers. Where there are benefits to the education and research community, Jisc may work with suppliers to make such services available, however this will always be done in ways that allow individual customer organisations a free choice between in-house, Jisc-provided and third-party provision.

As the Janet network community has grown (particularly with the inclusion of under-18s from schools and colleges) and with growing concern about the availability of unsuitable content on the internet, the issue of content filtering has become increasingly important. Even within an organisation, content filtering requirements will vary: material unsuitable for the younger pupils in a school may have educational value for their older peers. The approach taken by the Janet network allows each organisation or sector the flexibility to implement and manage appropriate content filtering for its particular requirements. The Online Safety Act 2023 is an example of legislation aimed at protecting users from harmful content online by imposing a range of duties on social media companies and search services, making them more responsible for their users' safety on their platforms.

Whilst there are a number of technical approaches to the challenge of preventing access to specific sites on the internet and filtering inappropriate material on the basis of its content, it should be recognised that these tools, either singly or in combination, will not be 100% effective in restricting access to inappropriate material. In addition, some types of filter may also block legitimate material that is either similar to, or located close to, the offending material. It is therefore recommended that any technical solutions deployed at a given site should be supported by other measures, including education on how to avoid and deal with inappropriate material and the implementation and enforcement of an Acceptable Use Policy (AUP) governing network access and computer use.

3 Acceptable Use Policy (AUP)

In common with many other network providers, Jisc defines an AUP, which can be found at: https://community.jisc.ac.uk/library/acceptable-use-policy [1]. It is a requirement for connection to the Janet network that each organisation agrees to comply with this Janet AUP. Having an AUP does not, and cannot, mean that non-compliant use will never occur. The purpose of the AUP is to ensure, as far as possible, that Janet services are used in an acceptable manner and in accordance with current legislation. Complying with the AUP also avoids the waste of resources, both for connected organisations and for the network, that is likely to be attendant on inappropriate or illegal usage. Connected organisations are expected to take reasonable measures to discourage behaviour that does not comply with the AUP and to take prompt and effective action to deal with complaints that the Policy has been breached. The Janet AUP does not require organisations to make it impossible for their users to act other than in accordance with it.

Connected organisations are strongly recommended to incorporate at least the same principles into their own local AUPs, along with any further site-specific requirements. They must also ensure that individual users are aware of the Janet AUP and the fact that they are bound by it. The consequences of failure to comply with AUPs (which may include suspension or revocation of computer and network access) should be made clear to users in advance. For avoidance of doubt, it may be helpful to have users sign an agreement to this effect when they are first granted network access. Reminders of the applicable policies — for example, through signs or login boxes — are often useful in maintaining the level of compliance.

4 Why Filter?

Filtering is often used to protect the networks and computers of connected organisations from hostile or unwanted network traffic. This type of filtering could be based either on the source of the traffic (such as an external site known to be sending large volumes of unsolicited bulk email) or on the destination (such as sensitive internal administration systems or subnets). Filtering unwanted large traffic flows can sometimes reduce network congestion, thus improving performance for priority applications, although this depends on the nature of the traffic and where it is blocked.

Filtering can also be used to force traffic to follow a particular route through the network. For example, an organisation may wish to implement a policy that all outgoing mail messages must go through an organisational mailserver, or that all web browsing must be done via the organisation’s web proxy/cache. The former policy can both limit the spread of viruses and

ensure that internal mail addresses are not revealed; the latter is essential to prevent users simply bypassing any content filtering performed by the proxy/cache.

Some organisations also wish to use filtering to prevent their users from accessing illegal or inappropriate material. It may be more challenging to implement filtering for this purpose than for those listed above, as many of the sites that contain inappropriate or illegal material will also contain material that is harmless or even beneficial. Simply blocking access to the whole site will exclude all of its content. To avoid this, a filter is needed that can determine the particular part of the site (often represented by a URL) that is being requested and decide whether to block or permit access at this more specific level. As mentioned above, the definition of appropriate may also vary between groups of users, so the settings for this type of filter may need to change frequently as different groups use the same computers.

Filtering policies are likely to need to allow for exceptions, both based on time and location. Technology and support processes must be designed to support these requests. For example a health education class might need to access selected websites about drugs that would otherwise be blocked; or a research group might legitimately be investigating illegal content. Guidance on the latter is available from Universities UK [2].

5 What Blocking or Filtering Does Jisc Provide?

The Janet network is designed and operated so as to allow connected organisations as much flexibility as possible in determining their most appropriate network service. At the network level, therefore, blocking is only used when required by the Janet Security Policy [3] as a limited measure to protect a particular organisation or service from an imminent technical security risk until the organisation has been able to address the problem itself. Because Jisc does not provide domain name resolution or content filtering by default, such blocking can only be imposed at packet/IP level, as discussed in 6.1 Packet filtering below.

Other than these temporary security measures, there is no centrally imposed filtering of web, email or other content provided by the network; indeed, such filtering would be ineffective as the network provides many possible routes to bypass any solution implemented at a single point.

Application-level services offered by Jisc may include filtering or blocking of some types of content, either as a permanent feature or as an option that can be enabled and configured by organisations using these services:

- The Janet Network Resolver Service [4] includes a frequently updated list of domain names that have been reported as compromised or otherwise malicious. When a user requests resolution of one of these domains, the service will instead return the address of a Jisc server that will attempt to inform the user of the threat. Organisations that subscribe to the JNRS must configure their local resolvers to use it, to benefit from this dynamic protection (see 6.2 Domain name filtering below).

6 Methods of Filtering Content

Since the Janet network does not contain built-in filtering of web or other content, connected organisations wishing to restrict access to specific material or sites on the Internet will need to

implement or obtain mechanisms to do this. They will also need to configure their own network routers or firewalls to block or re-direct deliberate or accidental attempts to bypass their filters. Filters that are managed by individual organisations (whether implemented at the organisation or elsewhere) and supported by local network and system configuration are both the most effective and least disruptive way to enforce the policy of the individual organisation.

Provided the local network is properly configured, the openness of the Janet network means that the filtering system can be equally effective wherever it is physically located. Appropriate filtering solutions may be available from local consortia, from Jisc (as described above) or from national or international providers. Four approaches to filtering, which may be used in combination, are in common use:

6.1 Packet filtering

Typically implemented on routers, the source addresses and port numbers of individual incoming IP packets are examined and compared against a banned list, and packets are only transmitted if there is no match. This approach results in blocking all traffic to or from the specific sites or networks in the banned list, or using a specified port number (which may correspond to a type of network service), whether the actual traffic is wanted or unwanted. If a blocked site uses a cloud host or content delivery network (CDN), packet filtering is likely to also block other customers of that host or CDN. The effort required to maintain the list of banned sites means that this approach is suitable only for fairly static lists that should be blocked for the whole organisation.

6.2 Domain name filtering

The Domain Name Service (DNS) is used to translate user-friendly names of internet sites (such as www.jisc.ac.uk [5]) into the numeric addresses used by computers to communicate over the network. Organisations can modify the behaviour of their DNS resolvers so that names associated with inappropriate content either do not resolve, or attempt to return an error page to the user. This can be done either on local resolvers, or by configuring computers to use one of a number of “safe DNS” services. Such blocks can be evaded by users changing their own DNS settings or entering the numeric addresses directly. Where a domain hosts a mix of wanted and unwanted content, a domain name filter will prevent access to both. As with packet filtering, domain name filtering is best suited to rules that apply to the whole organisation: creating exceptions for particular individuals, groups or times requires additional devices and configurations whose maintenance is likely to be complex and error-prone.

6.3 Application content filtering

This requires all off-site traffic to be routed through a proxy server which retrieves web pages on behalf of the requesting client system. The proxy server system runs software that can simply be configured to block access to entire sites based upon lists of banned addresses, as for packet filtering. However, proxy servers can also block access to specific web pages within a site by checking the web page address (or URL) or in some cases by examining the content of a requested page for specific keywords. This type of filtering can be more precise in the rules it applies, particularly where large websites contain only a minority of inappropriate material.

6.4 Hashlist filtering

Some types of illegal material are listed in catalogues of hash values. Individual images – for example indecent images of children – are examined, classified and added to these catalogues by organisations legally permitted to do what would otherwise be a criminal act. The use of hash values means that the images cannot be re-constructed from the catalogue, but that files with identical (or, in some cases, closely-related) content can be recognised. To use such a list, an organisation calculates the hash values of files stored on its own systems and compares these against the catalogue. Hash values are statistical, so a matching hash value is not conclusive proof that the file contents are the same, however such clashes are rare.

6.5 Implementing filtering

There are many commercial packages available which provide content filtering functionality, with regularly updated lists of banned sites that may be rated by category, age, or other factors. Reviews of a number of these packages are available through some of the web references given at the end of this document. The Squid project also offers a freeware package. It should be noted that Jisc does not recommend specific filtering software.

Different types of filtering lists are available, each of which is only effective if implemented at the appropriate layer of network technology:

- IP Address lists (e.g. 127.1.1.1) at routers or firewalls;
- Domain Name lists (e.g. www.example.com [6]) at DNS resolvers or firewalls;
- URL lists (e.g. http://www.example.com/unwanted-content.html [7]) in web proxies;
- Hash lists (e.g. of known illegal images) on filestores or content platforms.

The widespread use of cloud services, content delivery networks (CDNs) and dynamic DNS resolution makes it highly likely that using a list at the wrong level will result in both under- and over-blocking. For example converting a Domain Name to an IP address and blocking that address with a packet filter is likely to also block other domains that happen to use the same Internet host (and therefore share the address) and may well fail to block the desired domain if it is on a cloud service that distributes hosting across multiple machines.

These techniques are not and cannot be completely reliable for preventing deliberate or accidental access to inappropriate material. Lists of banned sites require regular maintenance, and so will not always be up to date. Additionally, there are well-known methods for evading the checks (e.g. the use of translation engines, or the embedding of redundant information in URLs). A further important consideration in the deployment of a proxy server is that it can introduce a potential point of failure into an organisation’s network infrastructure. If all network access is directed through a proxy server, then failure of that system can prevent all Internet access from client systems. Inadequate proxy server hardware can also result in (apparently) degraded network performance for users.

7 A Complementary Approach

It is suggested that organisations wishing to block access to internet content should adopt a multi-faceted approach to the problem, combining administrative, educational and technical

elements.

They should:

- Agree a policy about what internet content is suitable and what is unsuitable;
- Publicise that policy and incorporate its aims into an AUP;
- Ensure that all staff, students and visitors agree to comply with the AUP when first granted computer and network access, and make clear what the penalties are for non-compliance;
- Educate users in how to deal with inappropriate material they may find: in particular, encouraging them to report, rather than conceal, any accidental discovery of unsuitable material;
- Locate public access computers in open, supervised areas; if appropriate, requiring internet use to be accompanied or supervised;
- Implement technical measures where appropriate (for example, a proxy server) to enforce the policy on acceptable use. Such measures must be accompanied by appropriate configuration of the local network routers or firewalls, or they will be ineffective;
- Provide mechanisms to receive, assess and implement requests for temporary or permanent alterations to filtering rules;
- Use the monitoring capabilities of content blocking software to log network activity, and review the logs on a regular basis. Such monitoring must comply with the relevant legislation and users should be informed thattheir use will be monitored;
- Take appropriate action against any instances of non-compliance with the AUP.

8 Suggested Web Sites for Further Information

- NEN e-Safety site, covering safe use of the Internet for schools - https://www.nen.gov.uk/advice-for-schools/online-safety/ [8]
- Home page for the Internet Watch Foundation, a UK body concerned with the issue of illegal material on the Internet - ?http://www.iwf.org.uk/ [9]
- Universities UK guidance on research involving sensitive material - https://www.universitiesuk.ac.uk/what-we-do/policy-and-research/publications/oversight-security-sensitive-research [2]
- Janet Network Resolver Service - https://www.jisc.ac.uk/janet-network-resolver [4]
- Jisc Foundation GeoIP filtering - https://www.jisc.ac.uk/ddos-mitigation [10]

Annexe A How to think about Filtering

Different kinds of unwanted activity raise different issues around people, processes, and technologies. The following four questions (taken from a blog post [https://regulatorydevelopments.jiscinvolve.org/wp/2022/08/05/thinking-about-blocking/ [11]]) may be useful to explore whether and how filtering technology can contribute in specific cases:

* Where is the list?

* Where is the technology?

* Who are the users?

* How will people react?

Where is the list?

Any technology needs a set of instructions. In the case of blocking, we need to tell it how to distinguish things that should be blocked from things that should be allowed. Typically, that's a list of Internet locations. One day machine learning may get closer to understanding content or intention, but we'll still need to provide it with a good/bad model.

So, can we get that list from someone else, or do we have to create and maintain it ourselves? Maintained lists of different categories of activity may be available, either free or as part of commercial services or appliances. If we have to create a new list, do we have the skills, resources and permission (in some cases including legal) to do that? How will we keep it up to date, and handle any challenges to our decisions to include or exclude particular actions or content?

Where is the technology?

Internet technologies typically give us four different ways to specify things to be blocked: network (IP) addresses, domain names (DNS), application identifiers such as URLs and email addresses, and content inspection (e.g. keywords or hash values). Each of these gives a different precision, depending on the nature of the unwanted activity, so we should choose the one that most accurately defines what it is we want to block. Errors are likely in both directions – over-blocking that prevents legitimate activity: under-blocking that allows some unwanted – but choosing the right blocking mechanism should minimise these. Modern technologies such as cloud hosting and Content Delivery Networks (CDNs) involve a lot of sharing of both domain names and IP addresses, so those rarely offer good precision. Application identifiers are usually the most precise but extracting and checking them adds delay and privacy issues. Content inspection is unreliable outside a narrow set of applications, such as detecting repeat appearances of known illegal images.

Whatever technology layer we choose for blocking, we need some equipment to implement the block, and some way to ensure that network traffic goes through that equipment. Depending on the approach chosen, existing routers (IP), resolvers (DNS) or proxies (identifiers and content) may offer relevant functions: otherwise new equipment will be needed. Note that forcing traffic through blocking equipment is likely to create a single point of failure. Blocking and resilience are very hard to reconcile.

Who are the users?

A few kinds of activity – notably, active threats to connected computers – can be blocked for every user of the network. More often institutions will want to choose which blocks to apply and to whom, so should opt-in to the blocking, rather than having it imposed. If institutions need to make local changes to make blocking effective, imposing it before they are ready will have unpredictable results, possibly undermining existing protection measures. To assess the effectiveness of blocking, or to use the blocked content in research or teaching, particular individuals or locations will need to be exempted from the block.

These issues have implications for where the blocking equipment is located, who configures it and has access to logs. Equipment should be placed where it will have access to as much of the traffic to be checked as possible and (because most technologies add delay) as little other traffic as can be arranged. Where fine-grained per-user or per-location control is needed, this must be managed by the organisation that can identify the individuals and locations that should be (temporarily) exempted: typically their institution. Note that such fine-grained control is technically complex to implement for IP and DNS blocks. Where access to logs is required – for example to provide help to those who may have tried to undertake prohibited activities – this should also be at institutional level.

How will people react?

Technical blocks can always be circumvented, so are most effective against activity that no one should want to encounter. Even if recipients welcome the block, we still need to consider how malicious actors will respond: they might simply change location so we have to update lists more frequently; but they may also move activity closer to legitimate services to make over-blocking more likely and more disruptive.

Attempting to block activity that users desire gives them an incentive to circumvent the block. They can use different connectivity (home or mobile), but there are many technical ways to evade blocks without changing network. The activity may then continue but be invisible to those operating the network. Worse, most evasion technologies circumvent all blocks, including those for unwanted activity such as viruses, ransomware and other threats to devices and individuals. As our Guide to Filtering on Janet explains, it is particularly important that technical measures against desired content are part of a wider awareness, behaviour and support process: information and warnings may help reduce deliberate circumvention.

Examples

Two examples show how the questions can help explore the use of technology against different types of unwanted activity.

Distributed Denial of Service (DDoS)

- **Where is the list?** DDoS attacks against Janet and its customers are usually identified and blocked using a combination of source IP addresses and packet characteristics. Live information is available from commercial sources as well as Jisc’s own threat analysts.
- **Where is the technology?** Although some attacks can be blocked using existing routers it is more efficient (and less disruptive to the routers’ intended function) to

redirect suspect traffic to a dedicated cleaning service where malicious traffic can be identified and blocked and legitimate traffic from the same sources (which are usually compromised computers) forwarded to its intended direction.

- **Who are the users?** DDoS attacks can target any institution or service. Since blocks are typically temporary (the average attack duration in mid-2022 was one hour, the maximum six) and precise, they can be applied to protect all users of the Janet network.
- **How will people react?** Targets of blocked attacks should welcome the protection provided. Attackers can, and do, switch both the sources and targets of their attacks when blocked. Hence it is essential that blocks reflect live information from the network, as well as from external sources.

Terrorism

- **Where is the list?** Lists of content, including some regulated by Terrorism laws, are available through filtering services.
- **Where is the technology?** Terrorist content is frequently published through otherwise legitimate social media and hosting services. Unwanted content therefore needs to be defined at URL level, suitable for application proxies able to make these distinctions.
- **Who are the users?** UniversitiesUK has guidance on how to provide researchers with the access they need to security-sensitive material. Such access must be managed by the institution that can vet requests for access, verify the identities of authorised researchers, and provide appropriate access control facilities.
- **How will people react?** The Home Office Prevent Duty Guidance warns that some individuals may be drawn in by this kind of material. These may quickly adopt technologies to evade any blocks, so institutions' Prevent strategies should aim to provide appropriate advice and support to anyone showing early signs of interest.

Links
[1] https://community.jisc.ac.uk/library/acceptable-use-policy
[2] https://www.universitiesuk.ac.uk/what-we-do/policy-and-research/publications/oversight-security-sensitive-research
[3] https://community.jisc.ac.uk/library/janet-policies/security-policy
[4] https://www.jisc.ac.uk/janet-network-resolver
[5] http://www.jisc.ac.uk
[6] http://www.example.com/
[7] http://www.example.com/unwanted-content.html
[8] https://www.nen.gov.uk/advice-for-schools/online-safety/
[9] http://www.iwf.org.uk/
[10] https://www.jisc.ac.uk/ddos-mitigation
[11] https://eur01.safelinks.protection.outlook.com/?url=https%3A%2F%2Fregulatorydevelopments.jiscinvolve.org%2Fwp-about-blocking%2F&data=05%7C01%7CJohn.Chapman%40jisc.ac.uk%7C9e9b2f98772b4d04ade108da7b9a5e1b%