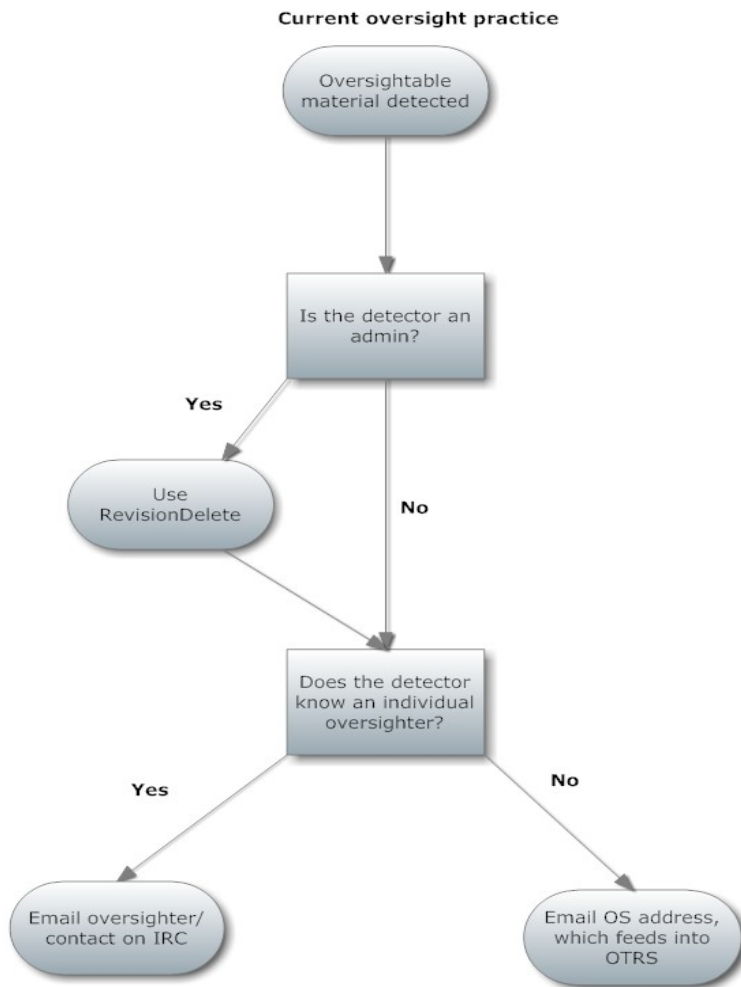


Oversight and the Article Feedback Tool

Current situation with oversightable content




At the moment, article content that is considered oversightable (particularly egregious copyright violations or personal attacks, private information etc) is detected, mostly, by normal users. If the user who detects a particular violation is an overseighter, the chain ends there; if they are an administrator, they first use the RevisionDelete function to hide it before reporting it. If they are any other user, they simply report it.

Reporting is done in one of two ways; if the detector knows an overseighter, they can simply email them or contact them on IRC (if they are available). Otherwise, an email is sent to the Oversight email address, either directly through an email client or by using the Special:EmailUser function for User:Oversight. Either way, the email goes to the "oversight" queue in OTRS (the Open Ticketing Response System, a customer support system used by our volunteers) which only overseighters have access to, where they then either oversight the content or leave it alone, depending on what they think about it.

To my knowledge, this has proven perfectly workable so far; I haven't heard any complaints from overseighters or requests for more support, although I will be contacting the Oversight mailing list to request input on whatever idea we settle on. Crucially the method isn't

particularly dependent on MediaWiki, which means we don't usually need to take it into account when making software tweaks – quite important when there are as many small fiddly components as there are in MediaWiki ;p.

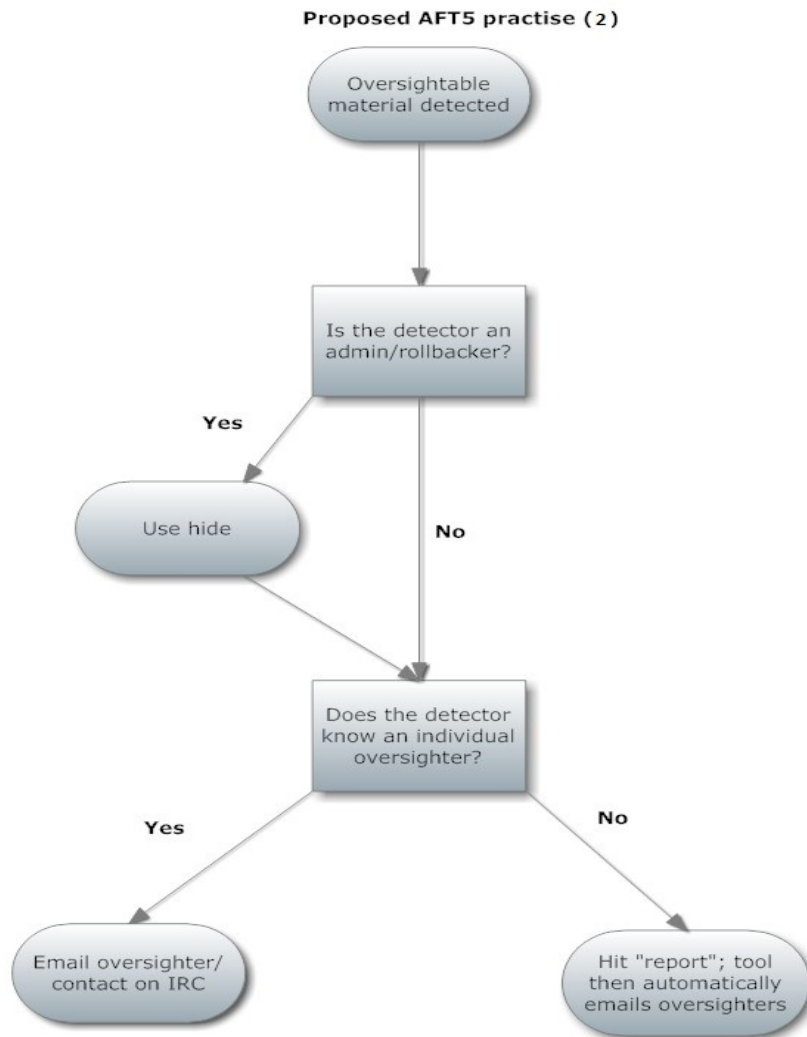
Proposed practise for AFT5 (option 1)



For AF5, option 1 is to essentially do the same thing as with article content. The advantage of this is that we avoid rocking the boat, and follow the principle of "if it aint broke, don't fix it"; since all comments will be accessible via individualised links, we can do precisely the same thing as is currently done with "diffs"; someone who notices inappropriate content brings it to an overseighter's attention via email or IRC, and the overseighter then deals with it.

The disadvantage of this is that we're designing new software which will require work from our volunteers. If we can do anything to make their lives easier, we should at least investigate it.

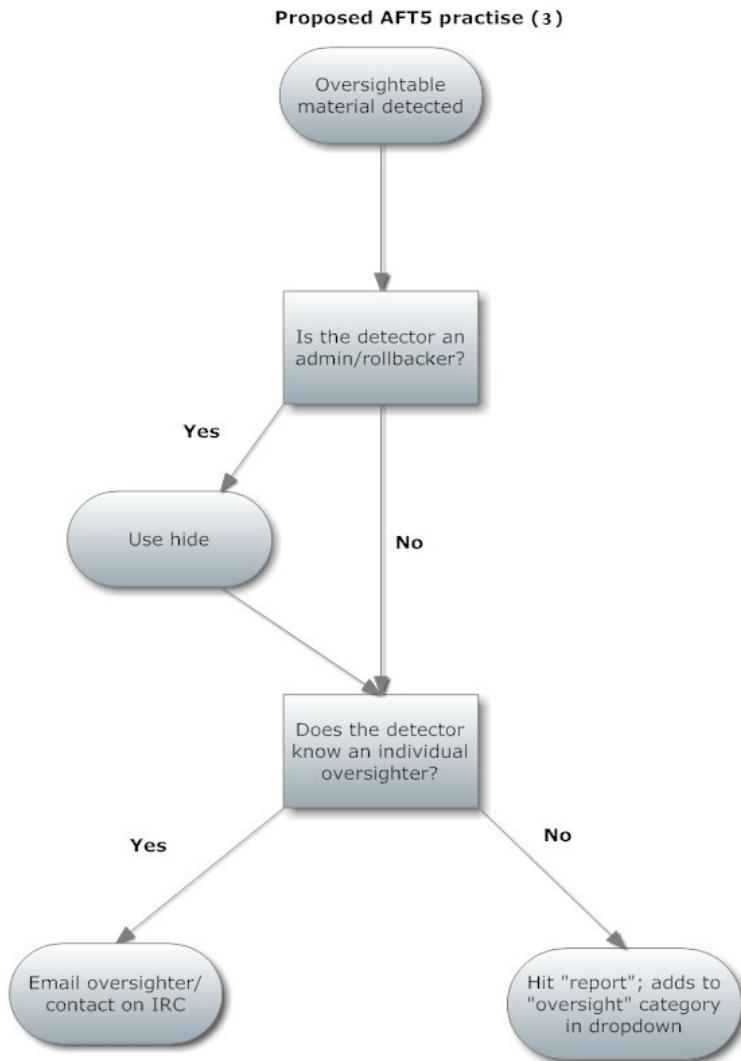
Proposed practise for AFT5 (option 2)



One idea to try and simplify things for the overseighters is to replicate the normal oversight process, but with a "report to overseighters" button next to each feedback item. This would automatically email the oversight address with a link to the specific piece of feedback, along with the username of the editor who has reported it. An overseighter could then follow the link and either confirm or deny oversight.

A problem with this is the potential for abuse; an unscrupulous user could merrily go along hitting the button and bombarding overseighters with deliberately cretinous requests for oversight. We can partially stop this by introducing the ability for overseighters to "deny" an oversight request – which also prevents the "report to overseighters" button being hit again for that particular piece of feedback – but that's only a temporary solution.

Proposed practise for AFT5 (option 3)



The third idea is pretty much identical to the second, but instead of sending an email, it simply adds the feedback to a special "oversight requested" category in the dropdown menu. To me this looks like the best way; it removes much of the risk with deliberate trolls or vandals (the oversight notification is *passive* rather than *active*, meaning improper requests bug the overseighters much less, reducing most of the incentive for pissing about), while still centralising requests. Whatever one we go for, if it involves a "report to overseighters" button it should definitely include, I think, the ability to "deny" a request and effectively turn off that button for that piece of feedback.