

Saving Wikipedia, Scaling Wikipedia

Copyright 2008 Greg London



This work is available under the Creative Commons-Attribution license

<http://creativecommons.org/licenses/by/3.0/us/>

With the following information provided for purposes of attribution:

Title: "Saving Wikipedia, Scaling Wikipedia"

Author: Greg London

URL: <http://www.greglondon.com/cc/by>

Version: 1 Jan 2008

Wikipedia is in the process of slowly imploding. Vandals and POV pushers are messing with the articles. Admins have to clean up the mess. But the job of administration requires a lot of oversight, which wikipedia does not have. Which makes selecting admins difficult and monitoring admins a difficult job. The problem is that while the vandals multiply easily and rapidly, the number of administrators does not scale easily to keep up.

The way Wikipedia is currently designed, the position of Administrator requires a lot of overhead for each person who is made an admin. The election process is rather lengthy, and once elected maintaining oversight on each admin invokes Brook's Law, increasing the amount of work that wikipedia has to deal with. As wikipedia continues to add admins, it could potentially reach a point where the total amount of effort put into creating and maintaining every administrator position and dealing with abusive editors, POV pushers, and so on, will become greater than the effort put into creating good encyclopedia articles.

Some might argue that this point has already been reached.

The problem is that the position of administrator does not scale well as the number of editors increase. The solution would be to make the application of administrative duties scale easily as the number of editors increase.

The solution is to redesign administrative duties so that these powers require relatively little overhead to hand out to people, so that there's little cost in adding admins to keep up with the increase in editors. The current role of administrator has a lot of overhead in the election process and a lot of overhead in the oversight process. If the roles were redesigned such that monitoring these admins was made easier by making it easier to spot abuse, by limiting the amount of potential abuse that any individual admin could potentially inflict, and by discouraging those individuals who would tend to be drawn to the role of administrator due to the power it gives them, then the overhead per admin would greatly drop as it would be easier to create new admins and monitor them for abuse. The number of admins could then easily scale to keep up with new editors coming to wikipedia.

This became the topic of conversation over at [Making Light](#). With the kicking around of some ideas, I think we came up with some workable solutions. This is a write up of one piece of it, dealing specifically with two new types of admins called “Janitors” and “Jurists” and a ticketing system for Jurists. There are some more suggestions on the Making Light thread regarding editor subtypes and some other ideas, but this document just got way too long to include everything.

First, there would be a new type of administrator called a “Janitor”. A Janitor can block users for two reasons: obvious vandalism and violations of a strict interpretation of the 3RR rule. Why these two violations? Because they are easy for people to quickly spot a violation, and they are easy to spot when a Janitor misuses their privileges.

Currently, the violations that Janitors would take care of are vandalism and 3RR. This could possibly be expanded to other violations, but it is very important that the job of Janitor be limited to rules that are obvious to see and obvious to detect misuse. This will allow scaling, which means a lot of Janitors can be easily added, and oversight will be minimal. Abuse will be obvious and the Janitor would have their privileges revoked.

Becoming a janitor should be relatively easy. It might not even need a nomination process or a vote for consensus. Which is also good, since these take a lot of time, and create an opportunity for quid pro quo among users. Instead, anyone with some minimum number of edits could become a janitor by request. Removing a janitor's privileges should be equally easy. An obvious misuse of privileges is an automatic and permanent revocation of privileges.

With an easily expandable group of Janitors dealing with Vandals and 3RR violations, the task left for the remaining Admins is one of dealing with POV pushing, violations of NPA policy, and other issues which generally require more of a subjective interpretation of events to see if a violation occurred.

This is more difficult to scale, because an admin can abuse their blocking privileges by applying a creative interpretation of what is a NPA violation, what is POV pushing, etc. Which means an abusive admin is harder to detect. The current solution is to fix this with a large up front cost of selecting admins through a nomination process, with the assumption that once nominated, the admin will operate without abusing their power.

However, this ignores the age old adage that power corrupts, and absolute power corrupts absolutely. While an admin may have the best of intentions at the time they are running for election, the current process for elections, with all of its massive overhead required to sort out good candidates, does absolutely nothing to deal with an admin who, with all manner of good intentions, begins to abuse their position and power after election. This situation is dealt with via the equally massive overhead process of taking an issue to the Arbitration Committee. This is why wikipedia doesn't scale easily.

What is needed is a way to redesign the position of Admin such that the possible ways of abuse are minimized, so that oversight isn't needed. The arbitration committee can still deal with abusive admins, but by redesigning the role of admins, it would be possible to reduce the potential for abuse in the first place, making it easier to add admins without worrying about abuse, and without adding massive oversight and overhead per admin.

The most immediate abuse is an admin who is editing an article and then uses admin privileges to block an editor who disagrees with their point of view. Wikipedia attempted to fix this by declaring that administrators cannot use their privileges on articles they are editing. An abusive admin easily circumvents this with allies who are admins. The admins work on their articles they are interested in, and when a dispute arises with a regular non-admin editor, the admin calls in their ally to block the editor using some creative interpretation of policy to find an excuse for a block. The admin repays this favor by going into the other admin's article upon request and returning the favor.

Note it's very important to distinguish the problem of “abuse” from the issue of “breaking the rules”. Currently, “abuse” isn't the same as “breaking the rules” until after a long, lengthy, tiring process before the arbitration committee. The idea here is to reduce to potential for “abuse” so that the overhead of per admin is reduced, and the job of admin can scale to keep up with editors.

Wikipedia's current system of rules easily and almost naturally encourages and creates an environment of quid pro quo among abusive admins. Even the process of nomination generally requires the support of existing admins. And that support may be the first “scratch” that the new admin feels indebted to repay.

This whole system of quid pro quo among the abusive admins is diffuse and subjective. It is diffuse because the guilty parties are spread out among several admins working as allies, and working slowly over time. And it is subjective because an abusive admin can use a subjective interpretation of events to declare some innocent behavior to be a violation and hand out a block. This diffusiveness and subjectiveness makes it extremely difficult to point to any one specific incident and say it was admin abuse. This in turn makes it extremely difficult to catch admin abuse. This difficulty makes any case before the Arbitration Committee extremely long and exhaustive. Which, in the end, makes it nearly impossible to scale the number of admins to meet the needs of wikipedia.

What wikipedia needs to be able to easily scale the number of admins is a way to short circuit the quid pro quo and the reward that comes from creative interpretation of rules. Requests for Deadminship are a lot of work, take a lot of time and energy, and simply do not scale to the number of admins needed.

The solution to this is another new type of admin called a “Jurist”. The title “Jurist” reflects the fact that the job of this particular admin is to judge the evidence, to bring in the subjective interpretation. A Jurist can block users for any rules violation, including NPOV and NPA, as well as vandalism and 3RR.

The difference, however, is that a Jurist cannot edit any content articles while they are a jurist. People who become Jurists must give up the ability to edit the content of any mainspace article. Jurists may edit talk pages, user pages, wikipedia policy pages, but not main content.

This breaks the first payoff that an abusive admin might use their privileges for. Editing an article and blocking anyone who gets in their way. While wikipedia's rule that no admin may use their privileges on a page they are editing, the response of an abusive admin is simply to build a group of allies and develop a system of quid pro quo so that an admin edits one page, and their friend hands out any blocks, and the admin returns the favor on the friend's page.

By disallowing the editing of main page articles, the set of rules make it much more difficult to create a system of allies and a system of quid pro quo with other jurists.

However, this still allows a Jurist to potentially misuse their privileges in two ways. First the jurist might develop ally editors. The editor works on the article sympathetic to the POV of the Jurist, and the Jurist hands out blocks to anyone who opposes that editor. Second the editor might develop a list of admins known to be friendly to his POV, and specifically ask for their help when dealing with the application of subjective rules.

To alleviate the natural tendency for allies to build between editors and Jurists, the system of Jurists will be handled with a Ticket Request system, or a docket system. A “request jurist” page will be created, and editors who witness another editor violating the rules will submit a request to the system.

When a Jurist has some time to work on a ticket, they will go to the Ticket Request system and ask for a case. The ticket request system will randomly assign a ticket to the Jurist. The Jurist can accept or decline the ticket. If they decline, that information is recorded publicly, and the Jurist is presented with another ticket. If the Jurist accepts the ticket, they then read the information given, render a judgment, and hand out a block.

This means that a Jurist cannot go into a page they are interested and hand out blocks to editors who oppose their POV. This also means that an editor cannot build a list of friendly Jurists and start a habit of making sure all requests for help aren't first directed to the friendly admin via email so that admin will be on the case. The ticketing system uses random assignments to prevent a jurist from selecting a page they have a particular personal POV about.

It might help to build into the system a record of all the users that a Jurist has to interact with on a ticket. And then any further requests that Jurist deals with should avoid those users. This should prevent any alliances from building.

So, Jurists have full admin privileges for blocking users for any policy violation. However, they are prohibited from editing article pages, and they can only operate on tickets handed out to them randomly via the ticket request system. This means that many of the incentives that encourage admin abuse have been removed from the system.

It also means that becoming a Jurist requires the user give something up.

Currently, becoming an Administrator grants the user many privileges without requiring the editor give anything up. This makes the job of administrator a coveted spot specifically for those who are inclined to abuse the position. While wikipedia states that being an admin is “no big deal”, the privileges handed out to admins with no privileges removed from the admin, mean that becoming an admin can only be seen as an increase in power. Whether that power is a “big deal” is a matter of each individual administrator, and oversight on that involves a lot of overhead.

By making Jurists give up the right to edit main page articles, those most likely to abuse an admin position are less inclined to apply for a Jurist position. Abuse is still possible, but it is limited to randomly selected incidents assigned to the Jurist. While a Jurist may have an intention to POV push on some political topic, the ticketing system would guarantee that they would only act as a Jurist on any particular page only once, making it hard for a Jurist to push a POV on a specific page.

All told, the job of Jurist then becomes something that can scale more easily than the job of Admin. Since most of the biggest incentives for potential abuse are removed, becoming a Jurist should attract people interested in contributing to a job that needs to be done on wikipedia, not people who want the power to assert their position on articles. This should make the nomination process much easier.

Being disallowed from editing main page articles also means that being a Jurist is more likely to be a temporary gig. Currently, being an admin on wikipedia gives an editor all the perks of being an editor with all the powers and privileges of being an admin. There is no incentive for admins to step down over time, and this can encourage extremely diffuse systems of quid pro quo, and admins extremely good at gaming the system to push their pov, as admins sit in power for several years.

If nothing else, the number of editors who become administrators, compared to the number of administrators who give up their admin powers and become regular editors again, should be sufficient evidence that being an admin is, in fact, a “big deal”. If it wasn't a big deal, then giving up the job of admin would be as common as taking on the job of admin.

A Jurist position might be more like a six month, nine month, or possibly one year gig. Editors give up the privilege to edit main page articles to volunteer for a job that wikipedia needs done. The Jurist serves their job for a while. And then the editor moves on and returns to editor status.

Wikipedia has a slogan that being an admin should be “No big deal”. By redesigning the system into Janitors and Jurists, this actually becomes a true statement.

The whole point of this idea of splitting the job of Wikipedia Administrator into Janitors and Jurists is to transform the job into something that can easily scale to keep up with the vandals and revert warriors and the POV pushers.

The job of Janitor is relatively easy for anyone to do, and relatively easy and quick to spot abuse, which both discourages abuse and makes spotting a bad Janitor and removing their privileges quick and easy too.

The job of Jurist is more complicated, but by removing the biggest incentives that would attract an abusive editor to want to become an admin, the Jurist system makes the abuse of admin privileges to push POV extremely difficult. Even if a Jurist wanted to abuse their position, the randomness of the ticket system limits the damage.

The intent is that by creating these two new positions, wikipedia can be saved from the onslaught of vandals and POV pushers without collapsing under the weight of editors looking to abuse admin privileges to push their own POV.

By scaling wikipedia, we can save wikipedia. By changing the job of admin to that of Janitors and Jurists, the role of policing wikipedia can be easily scaled, with little overhead, to keep up with ever growing onslaught of trolls and vandals.