# resilience



In order to reach its 2030 goals, the Wikimedia product platform must prepare for rapid scaling of development, of contributors and of content. In the process it will be critical to design for resilience – the ability to engender sustainable growth and to fend off threats. For example, it will be necessary to define countermeasures against external threats such as censorship, misinformation [1], climate and policy related threats, as well as attacks on security or privacy by state actors. It will also be necessary to anticipate and countermand threats that could undermine the projects from within: communities or affiliates turning against one another, communities turning against themselves [2] and communities turning against the Foundation. And finally, perhaps the most critical existential threat is relevance: what barriers to entry could be erected to prevent loss of mind share? What pre-emptive measures must be taken to guarantee mind share as new communities come on line? This paper explores each type of threat and offers a set of economic, cultural, and technical countermeasures. As the incumbent nonprofit internet presence defending a neutral point of view and access for all, it is critical that Wikimedia maintain and strengthen itself to preserve a future with truly free knowledge.

## Sections

# Scale 02

## Intro

Resilience refers to Wikimedia's ability to engender sustainable growth and to fend off threats. As the incumbent nonprofit internet presence defending a neutral point of view and access for all, it is critical that Wikimedia maintains and strengthens itself to preserve a future with truly free knowledge.

This paper recommends a number of countermeasures to be supported by the Wikimedia Foundation's Audiences department in order to bolster Wikimedia's resilience. A synthesized version on Officewiki is forthcoming as part of the Audiences department's 3-5 year planning FY 2018-2019.

## External Threats

There are four major external threats to Wikimedia:

- Censorship
- Misinformation, principally from state actors or sophisticated PR firms
- Climate- and policy-related disasters
- Attacks on security or privacy by state actors

## Internal Threats

There are also several ever present internal threats:

- Communities turning against each other
- Communities turning against newcomers
- Communities turning against themselves
- Communities turning against the Foundation
- The WMF turning against communities
- Wikimedia becoming irrelevant

## Countermeasures

The following alternatives seek to address a number of the threats listed above. There are no one-size-fits-all countermeasures for the threats, and thus a set of the alternatives would likely need to be applied for a robust defense.

**1. Consolidate to one domain name.** Consolidate Wikimedia production traffic under one domain name. This will discourage DNS poisoning and make DNS poisoning and TLS negotiation-based blocking more evident when it does occur.

**2. Give Huggle a hug.** Support growth and diversity of the editor ecosystem through targeted product enhancement: adapt (possibly mainstream) tools like Huggle with low BRD (Bold, Revert, Discuss) reciprocation rates. Make these tools run on additional contemporary platforms, adding features to streamline guidance to good faith editors, with integrated follow-up discussion, and promoting praise of edits going through this BRD cycle. Shepherd appropriately sized coalitions of users focused on the new platform tools and updated approaches.

**3. Decentralized internet distribution.** Work with key experts and OS and browser vendors to build a secure protocol stack for decentralized distribution that

- Ensures availability
- Maintains content integrity and recognizable URLs (e.g., Signed HTTP Exchanges)
- Shields reading habits from intermediaries (e.g., inbuilt browser tunneling or use of trusted peer nodes)
- Shields metrics logging from intermediaries (e.g., opaque out-of-band logging)
- Reasonably accommodates protecting readers from outdated reverted material for the common consumption case. [3]

This is in addition to other resilient Wikimedia hosted solutions. Forthcoming enhancements to core protocols (e.g., DNS over HTTPS and ESNI coupled with proxying through critical hosting intermediaries) present additional opportunities to raise the costs of eavesdropping and denial of service.

4**. Database copies to more cloud storage providers and mirrors.** More proactively place Wikimedia dumps on BitTorrent, Github, Gitlab, BitBucket, AWS/S3/Cloudfront, Azure, GCP, Rackspace, Akamai, and Cloudflare. Also foster more mirroring relationships with a global network of universities. Consider coordinating with Google, Cloudflare, and Bing to serve as hosts for AMP as a fallback of last resort in case of widespread system outage or blockade. Apply cryptographic signatures to these distributions. This would provide redundancy and would create obstacles to censorship while allowing experts to better verify edit histories.

**5. Two factor authentication.** Add support for two-factor authentication for all interested users. Holding all other factors constant (no pun intended!), this is one of the surest ways to confound a broad class of attacks on security and privacy.

**6. Invest in AI.** Consider further investment in AI resources for:

- Liar, outlier, and bias detection
- Machine vision and speech-to-text
- Labeling and model tuning

This will be necessary for combatting bad faith state actors and PR firms. It will also be necessary to support a probable influx of multimedia content that needs moderation (and tagging and translation). Product opportunities for high value micro-contributions abound here as well.

**7. Wikipedia all up.** Begin streaming of algorithmic or volunteer curator (or both) selected content via one or more of the following means. Consider a consolidated global Wikipedia brand. Offer language content in one to thirty languages, depending on the format.

- Internet radio
- Global radio frequencies
- YouTube (with permissive syndication)
- Multicast for broadcast and cable television
- Satellite TV

This helps in further establishing a global brand presence as an information utility, swinging the door open for further future investment. It also creates an outlet for Foundation and Movement thought leaders to explain how Wikipedia works and why. It also raises the costs of censorship at comparatively lower costs of support. Finally, it is an opportunity for forging collaborative user groups for durability and a global brand.

It should be noted this concept could easily be applied in native fashion on various consumer appliances as well, although that is a separate product question.

**8. Structured markup.** Embrace distribution on syndicating and interactive agent platforms, utilizing partnership conversations for bespoke treatments where appropriate. A broader presence not only keeps Wikimedia relevant, it makes suppression harder - for two reasons: (1) when Wikimedia is part of the fabric of life people won't take kindly to it disappearing, and (2) when Wikimedia is everywhere it's technologically harder to suppress. Employ five principles:

A. Use of structured markup. As specific next steps, (i.) add Schema.org support to TemplateWizard and (ii.) conduct a consultation with the Wikidata and major wiki communities about the Wikidata community modeling templates using Schema.org and weaving that modeling into the non-Wikidata projects (by mainstreaming of Parsoid markup). This is an opportunity to build trust between communities and help define some functional roles for the future.
B. Ability to measure impact. It's important to know if and to what extent distribution is helping the cause.
C. Clear branding. This is important for brand presence and enforcement.
D. Attribution. This is important for compliance and staying true to Wikimedia's values.
E. Positive contribution feedback loop. Not all distribution platforms will have this capability, but contribution should always be intentionally encouraged, and ideally, co-designed.

**9. Add Node.js and Python support to Templates.** Add support for Node.js, and possibly Python, to Scribunto. Scribunto supports the Lua language, which is not widely used. It should support Node.js, and possibly Python, which has a huge developer following.

Steer volunteer engineering toward

A. template (Scribunto) scripting, gadgets, and bots
B. improving MediaWiki Core

This places a higher emphasis on growing content and workflows for the wikis in a more sandboxed fashion while simultaneously making basic MediaWiki more excellent software for collaborative knowledge production (a global ecosystem form of resilience). Further investment in first class global templates, ideally with a mechanism to fuse data with Wikidata, is complementary. These new technologies are an opportunity to consider more contemporary code contribution workflows.

**10. Fund anti-surveillance and anti-censorship research.** Provide funding to 1-2 reputable anti-censorship / anti-surveillance firms (or fund incrementally internally). This lets more sophisticated forms of distribution and protection be developed.

## Summary

The following table is a guide to the countermeasures, how they address the major threats, their relative cost, and how the countermeasures might complement other efforts

| Countermeasure | Threats addressed | Cost and horizon | Complements |
|---|---|---|---|
| **1. Consolidate to one domain name**<br><br>Content and APIs are all served from wikipedia.org.<br><br>Censorship of one language is censorship of all, which is costly for censors. | → Censorship<br>→ Attacks on security or privacy by state actors | Medium, two year project with cross-functional team at 50% | Brand unification under Wikipedia |
| **2. Give Huggle a hug**<br><br>Get Huggle on Android and iOS. Improve its UX. Invest in productive in-app feedback loops.<br><br>People work nicely with each other, more editors stay around, the ecosystem flourishes. | → Communities turning against newcomers<br>→ Communities turning against the Foundation | Medium, three year project for one team | Making wiki projects bustling neighborhoods |
| **3. Decentralized internet distribution**<br><br>Host Wikipedia in a decentralized fashion with secure tunneling and digital signing.<br><br>Wikimedia is accessible even when servers are down or blocked. This is in addition to other resilient Wikimedia hosted solutions. | → Censorship<br>→ Climate- and policy-related disasters<br>→ Attacks on security or privacy by state actors<br>→ Wikimedia becoming irrelevant | Big, five year project for one small team with support from several other teams. Incremental milestones. | Eventually, offline editing |
| **4. Database copies to more cloud storage providers and mirrors**<br><br>Digitally verifiable database dumps become more pervasive. It's even harder to erase Wikimedia and its chain of edits. | → Censorship<br>→ Misinformation, principally from state actors or sophisticated PR firms | Small, one year project with one additional dedicated FTE | Academic research outreach |

| Countermeasure | Threats addressed | Cost and horizon | Complements |
|---|---|---|---|
| **5. Two factor authentication**<br><br>Anyone who wants it gets the option of two factor authentication. Account compromise becomes much harder. | → Attacks on security or privacy by state actors<br>→ The Foundation turning against communities | Medium, 18 month project with three dedicated FTEs and recurring SMS fees | Potentially, scoring and brand positioning |
| **6. AI**<br><br>Investment in liar, outlier, and bias detection;<br>machine vision and speech to text; labeling and model tuning.<br><br>It's easier to spot the bad guys. It's also easier and more fun for users to interact with and moderate content | → Misinformation, principally from state actors or sophisticated PR firms<br>→ Communities turning against each other<br>→ Communities turning against newcomers<br>→ Wikimedia becoming irrelevant | Large, 5 year project with paradigm shift for Audiences and Technology - varying levels of commitment by team. | Translation, scoring, mobile contribution and AI training, multimedia contribution, oral history |
| **7. Wikipedia All Up**<br>Content is streamed online, over the airwaves, and by satellite.<br><br>People can catch Wikimedia anywhere. Wikimedia is a trusted brand everyone knows will always be there, even for those without computers or smartphones. It's harder to block an omnichannel presence. | → Censorship<br>→ Misinformation, principally from state actors or sophisticated PR firms<br>→ Wikimedia becoming irrelevant | Medium, 3 year project with small team with escalating brand penetration | Brand unification under Wikipedia |
| **8. Structured markup**<br><br>Using structured markup and partner management, Wikimedia content is further embedded online, with impact measurement and Wikimedia values. Interactive agents automatically rely on the structured markup. | → Censorship<br>→ Wikimedia becoming irrelevant | Medium, 3 year project with small team with escalating brand penetration | Granulariz-ation of the article, translation |
| **9. Add Node.js and Python support to Templates**<br><br>Would-be template editors no longer need to use Lua for scripting (Scribunto), they can also use programming languages they know and love. This allows a key piece of the ecosystem to grow and thrive. | → Wikimedia becoming irrelevant | Medium, 2 year project with one dedicated FTE and one code review/tester | Global templates |
| **10. Fund anti-surveillance and anti-censorship research**<br>The next round of privacy tools gets researched and built while we pursue efforts on the current tools. | → Censorship<br>→ Attacks on security or privacy by state actors | 150K-300K annual investment | |

### Other Considerations

The following items are arguably not major Audiences efforts in and of themselves, but represent potential opportunities for other departments.

> Cooperate with Technology on a continuity plan in case both primary data centers go down for an extended period due to climate or policy disaster.
> Explore international governing body action on censorship on the basis of anti-competition (e.g., most blocks have corresponded with unfairly positioned state-supported alternatives) or adverse health and safety externalities (medical information and other critical information has become unavailable). This is a longshot, and the consequences for scrutiny on the content and the positioning as an NGO would need to be considered, but it may provide a defense.

### Conclusion & Other Opportunities

The countermeasures preempt future, and in some cases squash current, threats. You'll notice that they are also oriented around the space where the Wikimedia Foundation is uniquely positioned to take action, as these are large and difficult efforts requiring personnel. The recommendations don't fully capture the range of discussions or feedback received during late September and early October 2018 as part of the 3-5 year planning process.

Much more is said in other theme and subtheme documents as part of the Audiences 3-5 year planning point of view about potential community or feature interventions, but the following, which is heavily informed by recent conversations, are examples of how to aid in resilience in various other ways. Some of them overlap with material in other documents. They principally speak to creating the content and ecosystems that can activate and sustain growth, which is germane to the general theme of SCALE, as well as several other themes.

- Abuse filters
- Creating spaces to inform editors where there is surging demand or probable surging demand (based on algorithms) for topics and those topics do not yet meet a particular quality bar.
- Driving programs to encourage bilingualism. Exploring with professors the concept of translation proofreading as coursework.
- Ensuring inflows of translations into English Wikipedia and other major wikis.
- Investing in generalized work backlog solutions, catered for various personae and form factors.
- Emphasizing product experiences for mobile that are catered principally for AI training.
- Supporting federated SSO with major social identity providers, and flowing contribution activity back to user social channels.
- Scaling analysis of interventions by further integrating with academics in our data analysis.
- Partnering with a provider such as Coursera on a free course such as Programming Wikimedia: APIs, Bots, Gadgets, and Template Scripting.
- Supporting content snapshots (i.e., branded, perma-linkable, countable,

attributed hypermedia fragments) for embeddable content. This would be a complement to the summary endpoint and context cards.

## Notes

[1] Principally from state actors or sophisticated PR firms
[2] For example, veteran contributors working against newcomers.
[3] Note: risk concerning potentially infringing content, perhaps avoided by simply obfuscating discovery, needs analysis.

## Sources


A. Baso and O. Vasileva Research and Insights, Other contributors[1]: B. Davis, A. Halfaker, D. Horn, J. Katz, J. Minor, T. Negrin, M. Novotny, N. Pangarkar, S. Sastry, C. Ananian, A. Breault, C. Virtue, K. Thaney, J. Bennett, D. Foy, J. Vargas, R. Isler, B. Black, P.P.S. Naryan, SJ Klein

Early onset of structural inequality in the formation of collaborative knowledge, Wikipedia

Linguistic neighbourhoods: explaining cultural borders on Wikipedia through multilingual co-editing activity

Modeling crowdsourcing as collective problem solving

Robust clustering of languages across Wikipedia growth

The Rise and Decline of an Open Collaboration System: How Wikipedia's reaction to popularity is causing its decline

Chris Dixon essay Why Decentralization Matters, and selected sections of Build Your First Ethereup DApp.


[1] If your name was left off the list by mistake please contact JMinor or MNovotny