UvA-DARE (Digital Academic Repository)

Routes to Rights

Internet architecture and values in times of ossification and commercialization

ten Oever, N.; Beraldo, D.

DOI

[10.1145/3220561](https://doi.org/10.1145/3220561)

Publication date

2018

Document Version

Final published version

Published in

XRDS

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

ten Oever, N., & Beraldo, D. (2018). Routes to Rights: Internet architecture and values in times of ossification and commercialization. *XRDS*, 24(4), 28-31.
<https://doi.org/10.1145/3220561>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Routes to Rights:

Internet architecture and values in times of ossification and commercialization

The commercialization and evolution of the internet have affected our ability to exercise human rights online.

By *Niels ten Oever and Davide Beraldo*

DOI: 10.1145/3220561

The complexity of the internet is ever increasing: Year by year, new layers of interrelated technologies are sedimenting over existing ones. Certainly, this has allowed for great progress in terms of functionality and performance. However, this stratification process comes along with apparent downsides: the ossification of the infrastructure, centralized ownership, the rise of proprietary middleboxes, and decreased transparency. In this article we argue certain major negative features, inherent to the incremental development model of the internet, are not sufficiently taken into consideration in the design of new internet protocols. More particularly, greater emphasis should be put on the impact of protocols and infrastructures on our lives. Human rights—as laid out in

the Universal Declaration of Human Right, the International Covenant of Civil and Political Rights, the International Covenant on Economic Social and Cultural Rights, and other related treaties—are the most universal values we have, as almost every country in the world has ratified them. Given the internet’s universalistic vocation, we should ensure it does not harm people’s ability to exercise their human rights. We suggest a necessary move in this direction is to leverage the “right to science,” in order to open up internet architecture and infrastructure, re-involving the academic community in their design and maintenance.

OSSIFICATION OF THE INFRASTRUCTURE

Infrastructures notoriously present a high level of inertia: Once they are set, it can be very hard to change them. This means innovations need to take into account the existing architecture and cope with the issue of backward compatibility. In hindsight, it is evident the key technology of an information society is built upon a kaleidoscopic stack of interventions that resemble workarounds, convenient hacks, and quick fixes. This model of “permissionless innovation,” similar to the bazaar model of open-source software development, has allowed a

great level of intellectual freedom and creativity. Nonetheless, this inherent lack of purposeful architecture also leads to increasing ossification and complexity. The fact that this approach has never evolved has less to do with substantial reasons (related to technological rigor or to the public good) than with economic and political interests (such as the existence of quasi-monopolies and the nature of decision making within internet governance bodies). The internet governance ecosystem is built on a rather delicate equilibrium; thus, revising previous decisions can be quite a disruptive element. Instead, the strategy, so far, has

been to create updates and ensure new technologies are backward compatible with legacy ones. The consequences of this approach are troubling; new layers are regularly added to patch problems of existing layers, without ever fixing the actual problem. This ossification process means new solutions are built upon and depend on old mistakes.

We can mention two well-known examples of this vicious circle of workarounds that, in turn, reinforce existing problems: tunneling and network overlays; and recent calls for “network self-management” through machine learning algorithms. Tunneling and network overlays allow for groups or

organizations to virtualize part of their communication, thus making it less visible to third parties while relying on a common communication infrastructure. The downside side of this fix, however, is such a network adds a new layer that can, potentially, introduce new vulnerabilities. The calls for the use of machine learning to engage in network management might result in more efficiency and higher throughput. However, it makes the process of network management even less transparent than it is today. It would add a whole new category of middleboxes to the network with the ability not only to route packets, but

also to drop packets they do not recognize, thus hampering permissionless innovation. An example of this scenario recently emerged in the standardization of Transport Layer Security (TLS) 1.3, when the new security protocol could not be deployed as it was designed because middleboxes would not recognize it and, consequently, would drop the packets. In order to solve this issue, TLS1.3 was made to look like TLS1.2 so that it could “fool” the middleboxes. In other words, it seems in order to progress, we are doomed to mimic the past. Again, this happens not because of rational design decisions, but simply because ex-

isting non-standardized implementations are not made future-proof.

CENTRALIZATION OF OWNERSHIP

In the earlier days of the internet, its backbone was run by the National Science Foundation Network (NSFNET), which restricted commercial use of the internet through its acceptable use policy. This changed with the introduction of the Commercial Internet Exchange, and other internet exchanges that followed, leading to the decommissioning of the NSFNET and to the decentralization of ownership structures.

Just a couple of decades afterwards, internet ownership structure is now notoriously centralized. This does not only take into account the most visible components owned by Amazon, Google, Facebook, Apple, and Microsoft; even lower layers of the internet face the same trend. Giant industry players include companies in the fields of content distribution networks (with actors such as Akamai and Cloudflare), telecommunication provision (such as China Mobile, Verizon, AT&T, and Vodafone), networking hardware (such as Cisco and Huawei), and chip manufacturing (like Intel and AMD). These market players have a tremendous influence on internet infrastructure, but their code and products are proprietary and, thus, invisible—at least until they fail. We have seen this happen recently, when Amazon Web Services broke down and when serious security issues were detected in Intel chips. Of course, these incidents could have also occurred within a less concentrated market, but their impact would have been much smaller, and alternatives would have been available.

Another issue with market concentration is the presence of monopolies or oligopolies. There is much less motivation to set common standards among competitors, which in turn reinforces the market position of the monopolist and undermines one of the founding principle of the internet: interoperability. This trend also coexists with the increasing importance of network management devices known as middleboxes, which are breaking another of the internet's founding value: the end-to-end principle. Whereas the rhetoric of the internet as a liber-

The key technology of an information society is built upon a kaleidoscopic stack of interventions that resemble workarounds, convenient hacks, and quick fixes.

ating and empowering technology is still alive, the more the internet is becoming ubiquitous and integrated into the woodwork of society, the more its control is concentrated in the control rooms of a limited number of powerful actors. The original promise and premise of the end-to-end principle is compromised by the nefarious impact middleboxes have on the network. The concentration of ownership does not only mean reduced agency; it also makes the impact of vulnerabilities or network outage much larger.

To be sure, it has to be said this scenario seemingly has some advantages. The development of the Quick UDP Internet Connections (QUIC) protocol, for instance, is primarily driven by Google, which can easily lead the process because it is deploying QUIC in the communication between the Chrome browser and Google servers. QUIC improves privacy, security, and authentication, as well as performance through congestion control. But the reason why it can be adopted this quickly is because Google owns a large part of the market. This means privacy might be protected against third parties, but the “benevolent dictator” represented by Google could still harvest private data. If we compare this example to the uptake of Internet Protocol version 6 (IPv6), Border Gateway Protocol Security (BGPSEC), Resource Public Key Infrastructure (RPKI), Domain Name System Security Extensions (DNSSEC), DNS-based Authentication of Named Entities (DANE), and Domain Keys Identified Mail (DKIM),

it is clear how the QUIC case could be taken as a role model. However, if we take a closer look at its impact, we cannot help but recognize performance and security largely improved between Google browsers and Google servers, thus reinforcing Google's predominant position. The problem with alternative solutions is they only work for autonomous systems, or providers that already implement them, thus facing a notorious problem of low adoption rates. We can clearly see this in the case of the deployments of BGPSEC, RPKI, DNSSEC, DANE, and DKIM.

TRAGEDY OF THE COMMONS OR BENEVOLENT DICTATORSHIP?

While power is concentrated in the hands of a limited number of gigantic corporate actors, users and policy makers are left with a dilemma: Is it better to have a few monopolistic companies that can be held to account, or more diversity through smaller companies that might not, however, have the capacity (or priority) to live up to high standards? To be sure, the research and deployments done by large companies have come with advantages: bandwidth has grown larger, the availability of internet has increased, and latency is getting lower. On the other hand, transparency of routes, anonymity, data portability, and privacy are all decreasing. Moreover, internet connectivity prices have not yet dropped despite infrastructural improvements. The reasons for such a scenario are not hard to guess if we consider the economic interest of these companies, but what are the technical consequences of this observation? In order to answer this question, it might be a good idea to ask the technologists.

For too long the internet's technical bodies have acted as if their role was merely to facilitate the market and to make politically neutral technical choices. It is important to stress this development has coincided with the exodus of academics from technical standard-setting bodies, the virtual absence of a global civil society in many of these bodies, and the limited mandate and moral responsibility these organizations assign to themselves. This alignment of trends might not be a coincidence, since the leadership of these organizations consists of the same

companies for which these standards are being developed. The internet, in other words, has become a public good based on a private infrastructure, where a limited number of private companies design what everybody's future will look like. Nonetheless, we are somehow still clinging to the dream that the internet is an anarchic, distributed architecture controlled by everyone and no one. We might be too afraid to wake up and realize we are trapped in a big shopping mall.

It has to be said, there have been cases in which the actions of technical bodies dominated by big corporations have naturally aligned with the interests of end users. For instance, the Snowden revelations led to the Montevideo statement in which a number of technical bodies denounced mass-surveillance; however, this can easily be attributed to the companies' concern with compromising the trust in their products, and thus profit. Can end-user values (such as internationalization, affordable pricing, anonymity provision, data portability, etc.) be promoted within the current governance model? None of these values, unfortunately, seem to be in the direct economic interest of today's guardians of the internet.

Protocols that would cope with at least one of the many issues of today's internet (IPv6, BGPSEC, DNSSEC, DANE, and DKIM are all related to security) are facing problems of low adoption, which the standards bodies do not seem able to address. Internet infrastructure, in other words, is facing a collective action problem: Nominally, everybody has an interest in aligning with end-users' values; practically, nobody is willing to invest in it because of the cost associated with initial low market adoption. It would be great if companies within standard bodies could spontaneously agree on reforming the architecture based on end-users' values. However, if we have to rely on market mechanisms only, the best we can get is, from time to time, the "benevolent" concession of a quasi-monopolistic corporation with its agenda attached to it.

REINVENTING THE INVENTORS

It is an open secret that internet technical bodies are facing a problem of decreased market interest, which

translates to less participation. As a consequence, more and more developments take place outside of these bodies, and tend to focus on the application layer. There are few ways in which we can interpret this trend: 1) the working methods and timelines of these technical bodies do not fit current practices; 2) the infrastructure is done, people are moving; or 3) the problems are not interesting or relevant enough.

Whereas timelines can be indeed quite tedious, the added value of this timing is also significant. The answer to the previous question, then, more realistically lies in the lack of interest from market parties to engage in further standardization: They have developed their earning model and do not see much potential for increasing their bottom line through the standardization process. This does not mean the infrastructure "is done," but simply it is harder to make more money from it.

There is a clear need of revitalization; thus, this is exactly the good moment to evaluate what went wrong in the development of the infrastructure and ask how it could be re-imagined from the vantage point we reached today. This should finally include, we argue, the understanding that the internet mediates the ability of people to exercise their human rights. This is not only confined to more obvious aspects, such as security and privacy, but also to others, such as the right to freedom of expression, education, association, non-discrimination, equal protection, political participation, and last, but not least, the right to participate in cultural life, arts, and science.

Article 27 of the Universal Declaration on Human Rights and Article 15 of the International Covenant on Economic, Social and Cultural Rights, jointly constituting the basis for the right to science, recognize everyone has the right to enjoy the benefits of scientific progress and its applications. In order to uphold this right, states need to respect the freedom to pursue scientific research and take measures to develop and disseminate science. There is also an increasing trend, most notably through the evocation of the United Nations Guiding Principles for Busi-

ness and Human Rights, to keep businesses accountable for their impact on human rights. Under this responsibility and agenda one could seek to open internet infrastructure and re-involve the most notably absent from this discussion: the academic community, which was so crucial in the birth of internet architecture and now so silent in its current stage.

Why would such a scenario be relevant for the business sector? There is a global trend of discussions on cybersecurity taking place outside of technical bodies, with decisions being made that might not be consistent with the technical rigor of the network architecture. Engineers should be encouraged and supported to switch from being mere profit-enablers to becoming true inventors again. This would also be an homage to early internet developers and would allow us to build a sustainable, rights-respecting future based on a public-private partnership. Eventually leading to a significant impact on trademarks, patents, APIs, and data portability, and thus counter the process of knowledge privatization while reinvigorating the public debate on what the internet (and, consequently, the world at large) should look like.

Acknowledgment

This project has partly received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme [grant agreement No 639379-DATACTIVE].

Biographies

Niels ten Dever is a Ph.D. candidate with the DATACTIVE Research Group at the Media Studies department at the University of Amsterdam. His research focuses on the evolution of the notion of public interest in internet architecture. His other research interests include global governance innovation and the socio-technical ordering that invisible infrastructures provide to our societies. Previously, ten Dever has worked as Head of Digital for the freedom of expression organization ARTICLE19, where he set up the digital program that covered human rights and the IETF, ICANN, IEEE, and ITU. Before that, he designed and implemented international freedom of expression projects with Free Press Unlimited.

Davide Beraldo is a postdoctoral researcher at the University of Amsterdam (Department of Media Studies), where he is part of the DATACTIVE group, researching on the politics of datafication. He obtained his Ph.D. [*cum laude*] in social sciences and sociology, as a joint degree between the University of Amsterdam and the University of Milan. He holds a bachelor's and a master's in political sciences and social sciences from the University of Milan, and has a background in computer science. His research interests include social movements, digital media, algorithms, computational social sciences, and epistemology of complexity.

© 2018 Copyright held by author.
Publication rights licensed to ACM
1528-4972/18/06 \$15.00