

Principles for Open Scholarly Infrastructures

infrastructure [ˈɪnfərˌstrʌktʃər] (noun) – the basic physical and organizational structures and facilities (e.g. buildings, roads, power supplies) needed for the operation of a society or enterprise. – *New Oxford American Dictionary*

Everything we have gained by opening content and data will be under threat if we allow the enclosure of scholarly infrastructures. We propose a set of principles by which Open Infrastructures to support the research community could be run and sustained. – Geoffrey Bilder, Jennifer Lin, Cameron Neylon

Cite as “Bilder G, Lin J, Neylon C (2015) *Principles for Open Scholarly Infrastructure-v1*, retrieved [date], <http://dx.doi.org/10.6084/m9.figshare.1314859>”

Over the past decade, we have made real progress to further ensure the availability of data that supports research claims. This work is far from complete. We believe that data about the research process itself deserves exactly the same level of respect and care. The scholarly community does not own or control most of this information. For example, we could have built or taken on the infrastructure to collect bibliographic data and citations but that task was left to private enterprise. Similarly, today the metadata generated in scholarly online discussions are increasingly held by private enterprises. They do not answer to any community board. They have no obligations to continue to provide services at their current rates, particularly when that rate is zero.

We do not contest the strengths of private enterprise: innovation and customer focus. There is a lot of exciting innovation in this space, much it coming from private, for profit interests, or public-private partnerships. Even publicly funded projects are under substantial pressures to show revenue opportunities. We believe we risk repeating the mistakes of the past, where a lack of community engagement lead to a lack of community control, and the locking up of community resources. In particular our view is that the underlying data that is generated by the actions of the research community should be a community resource – supporting informed decision making for the community as well as providing as base for private enterprise to provide value added services.

What should a shared infrastructure look like? Infrastructure at its best is invisible. We tend to only notice it when it fails. If successful, it is stable and sustainable. Above all, it is trusted and relied on by the broad community it serves. Trust must run strongly across each of the following areas: running the infrastructure (governance), funding it (sustainability), and preserving community ownership of it (insurance). In this spirit, we have drafted a set of design principles we think could support the creation of successful shared infrastructures.

Governance

If an infrastructure is successful and becomes critical to the community, we need to ensure it is not co-opted by particular interest groups. Similarly, we need to ensure that any organisation does not confuse serving itself with serving its stakeholders. How do we ensure that the system is run “humbly”, that it recognises it doesn’t have a right to exist beyond the support it provides for the community and that it plans accordingly? How do we ensure that the system remains responsive to the changing needs of the community?

- **Coverage across the research enterprise** – it is increasingly clear that research transcends disciplines, geography, institutions and stakeholders. The infrastructure that supports it needs to do the same.
- **Stakeholder Governed** – a board-governed organisation drawn from the stakeholder community builds more confidence that the organisation will take decisions driven by community consensus and consideration of different interests.
- **Non-discriminatory membership** – we see the best option as an “opt-in” approach with a principle of non-discrimination where any stakeholder group may express an interest and should be welcome. The process of representation in day to day governance must also be inclusive with governance that reflects the demographics of the membership.
- **Transparent operations** – achieving trust in the selection of representatives to governance groups will be best achieved through transparent processes and operations in general (within the constraints of privacy laws).
- **Cannot lobby** – the community, not infrastructure organizations, should collectively drive regulatory change. An infrastructure organisation’s role is to provide a base for others to work on and should depend on its community to support the creation of a legislative environment that affects it.
- **Living will** – a powerful way to create trust is to publicly describe a plan addressing the condition under which an organisation would be wound down, how this would happen, and how any ongoing assets could be archived and preserved when passed to a successor organisation. Any such organisation would need to honour this same set of principles.
- **Formal incentives to fulfil mission & wind-down** – infrastructures exist for a specific purpose and that purpose can be radically simplified or even rendered unnecessary by technological or social change. If it is possible the organisation (and staff) should have direct incentives to deliver on the mission and wind down.

Sustainability

Financial sustainability is a key element of creating trust. “Trust” often elides multiple elements: intentions, resources and checks and balances. An organisation that is both well meaning and has the right expertise will still not be trusted if it does not have sustainable resources to execute its mission. How do we ensure that an organisation has the resources to meet its obligations?

- **Time-limited funds are used only for time-limited activities** – day to day operations should be supported by day to day sustainable revenue sources. Grant dependency for funding operations makes them fragile and more easily distracted from building core infrastructure.
- **Goal to generate surplus** – organisations which define sustainability based merely on recovering costs are brittle and stagnant. It is not enough to merely survive it has to be able to adapt and change. To weather economic, social and technological volatility, they need financial resources beyond immediate operating costs.
- **Goal to create contingency fund to support operations for 12 months** – a high priority should be generating a contingency fund that can support a complete, orderly wind down (12 months in most cases). This fund should be separate from those allocated to covering operating risk and investment in development.
- **Mission-consistent revenue generation** - potential revenue sources should be considered for consistency with the organisational mission and not run counter to the aims of the organisation. For instance...
- **Revenue based on services, not data** – data related to the running of the research enterprise should be a community property. Appropriate revenue sources might include value-added services, consulting, API Service Level Agreements or membership fees.

Insurance

Even with the best possible governance structures, critical infrastructure can still be co-opted by a subset of stakeholders or simply drift away from the needs of the community. Long term trust requires the community to believe it retains control.

Here we can learn from Open Source practices. To ensure that the community can take control if necessary, the infrastructure must be “forkable.” The community could replicate the entire system if the organisation loses the support of stakeholders, despite all established checks and balances. Each crucial part then must be legally and technically capable of replication, including software systems and data.

Forking carries a high cost, and in practice this would always remain challenging. But the ability of the community to recreate the infrastructure will create confidence in the system. The possibility of forking prompts all players to work well together, spurring a virtuous cycle. Acts that reduce the feasibility of forking then are strong signals that concerns should be raised.

The following principles should ensure that, as a whole, the organisation *in extremis* is forkable:

- **Open source** – All software required to run the infrastructure should be available under an open source license. This does not include other software that may be involved with running the organisation.

- **Open data (within constraints of privacy laws)** – For an infrastructure to be forked it will be necessary to replicate all relevant data. The CC0 waiver is best practice in making data legally available. Privacy and data protection laws will limit the extent to which this is possible.
- **Available data (within constraints of privacy laws)** – It is not enough that the data be made “open” if there is not a practical way to actually obtain it. Underlying data should be made easily available via periodic data dumps.
- **Patent non-assertion** – The organisation should commit to a patent non-assertion covenant. The organisation may obtain patents to protect its own operations, but not use them to prevent the community from replicating the infrastructure.

Implementation

Principles are all very well but it all boils down to how they are implemented. What would an organisation actually look like if run on these principles? Currently, the most obvious business model is a board-governed, not-for-profit membership organisation, but other models should be explored. The process by which a governing group is established and refreshed would need careful consideration and community engagement. As would appropriate revenue models and options for implementing a living will.

Many of the consequences of these principles are obvious. One which is less obvious is that the need for forkability implies centralization of control. We often reflexively argue for federation in situations like this because a single centralised point of failure is dangerous. But in our experience federation begets centralisation. The web is federated, yet a small number of companies (e.g., Google, Facebook, Amazon) control discoverability; the published literature is federated yet two organisations control the citation graph (Thomson Reuters and Elsevier via Scopus). In these cases, federation did not prevent centralisation and control. And historically, this has occurred outside of stewardship to the community. For example, Google Scholar is a widely used infrastructure service with no responsibility to the community. Its revenue model and sustainability are opaque.

Centralization can be hugely advantageous though – a single point of failure can also mean there is a single point for repair. If we tackle the question of trust head on instead of using federation as a way to avoid the question of who can be trusted, we should not need to federate for merely political reasons. We will be able to build accountable and trusted organisations that manage this centralization responsibly.

Is there any existing infrastructure organisation that satisfies our principles? ORCID probably comes the closest, which is not a surprise as our conversation and these principles had their genesis in the community concerns and discussions that led to its creation. The ORCID principles represented the first attempt to address the issue of community trust which have developed in our conversations since to include additional issues. Other instructive examples that provide direction include Wikimedia Foundation and CERN.

Ultimately the question we are trying to resolve is how do we build organizations that communities trust and rely on to deliver critical infrastructures. Too often in the past we have used technical approaches, such as federation, to combat the fear that a system can be co-opted or controlled by unaccountable parties. Instead we need to consider how the community can create accountable and trustworthy organisations. Trust is built on three pillars: good governance (and therefore good intentions), capacity and resources (sustainability), and believable insurance mechanisms for when something goes wrong. These principles are an attempt to set out how these three pillars can be consistently addressed.

The challenge of course lies in implementation. We have not addressed the question of how the community can determine when a service has become important enough to be regarded as infrastructure nor how to transition such a service to community governance. If we can answer that question the community must take the responsibility to make that decision. We therefore solicit your critique and comments on this draft list of principles. We hope to provoke discussion across the scholarly ecosystem from researchers to publishers, funders, research institutions and technology providers and will follow up with a further series of posts where we explore these principles in more detail.

The authors are writing in a personal capacity. None of the above should be taken as the view or position of any of our respective employers or other organisations.

[Top](#)