

The Hydra Way (draft text)

This was a version 0.1 which has been replaced by [The Samvera Way: a recipe for success](#)

There are those, we are told, who point to the Hydra project as one of the most successful, current, open source projects. Whether we deserve that accolade is up to others to judge but we in the Hydra Community **are** proud of what we have done and what we do. These few pages are an attempt to record just what was and is "The Hydra Way", and why we think it was and is successful.

In the beginning

The Hydra story begins in summer 2008 at the Open Repositories conference in Southampton, UK. After a presentation given by Richard Green on the early repository work being done at the University of Hull, Thorny (Thornton) Staples - then on the staff with the newly formed formal Fedora Commons organization - told him how Hull's work was very close to something that his recent colleagues at UVa wanted to do and that it seemed to him there might be benefit in collaborating. Thus it was that Richard and his colleague Chris Awre agreed to go to UVa in Charlottesville that September to discuss such a project. During the intervening months, Tom Cramer at Stanford came up with some similar needs and so represented around the table at the September meeting were Hull, Virginia, Stanford and Fedora Commons.

From the start this was a true collaboration: there was no senior partner and the meetings (of which there were many over the next twelve months) did not even have a formal Chair. Around the table all were equal and everyone's opinions were valued: a practice that we continue to this day.

Over the next twelve months or so, the partners met roughly every six weeks. This represented a huge commitment of resource, both human and financial. Hydra had no source of funding; the universities involved pursued the project because they believed in its goals and in the long-term benefits of a collaboration. Somewhere in those months we adopted the phrase "if you want to go fast, go alone. If you want to go far, go together", apparently an African proverb.

The universities weren't far into that series of meetings before the partnership was opened up to a fifth player, MediaShelf LLC. MediaShelf's founder, Matt Zumwalt, was very interested in the work and quickly became deeply involved. Some short-lived, initial disquiet about the involvement of a "commercial company" in the project was quickly dispelled as Matt demonstrated his total commitment to the project goals.

So what were the aims of Hydra? Essentially it was a collaboration to combine the partners' individual repository development efforts into a collective solution with breadth and depth that exceeded the capacity of any of the institutions to create, maintain or enhance on its own. But actually it was more than that. There was a realization, almost from day one, that a four- or five-way collaboration might be good but that a six- or seven- or 20-way one might be better. It was always the intention of the original team that more collaborators would be brought into the mix as time went on and that, in any case, all the work done for the project would be released as open source code in the hope that others would want to adopt it.

There was never going to be "a Hydra product" in the sense of "here is the DVD, go away and install it." Hydra was conceived as a set of building blocks that could be assembled in different ways at different institutions so that each Hydra solution (or Hydra head, as they became known) was uniquely tailored to local needs. The term "Lego bricks" was bandied around an awful lot in the early days. The original team were already committed to the use of the Fedora repository software as the underlying storage architecture and 30 minutes of demonstration from, what was then, UVa's Blacklight team convinced us that Blacklight had immense possibilities for driving Hydra's front end. Given its underlying technologies the adoption of Solr and Ruby on Rails (RoR) followed logically. Hydra, now the Hydra gem set, became the orchestration between these four underpinnings.

Are there lessons here for the "Hydra Way" here? For sure.

The first might be "trust people". Potentially there were a lot of egos around that first meeting table and yet a true collaboration of equals emerged. What we did come to realise was important was that team need to meet in person regularly in order to maintain the highest level of trust. Air fares or no, Skype conferences are no long-term substitute for meetings and a bit of face-to-face social time periodically. In practice, we try now to get the Hydra community together three or four times each year.

A second might be "seize opportunities". The adoption of Blacklight (and hence Solr and RoR) was not planned. It was an opportunity that came along, looked very promising and we grabbed at it. It's a grab we have never regretted.

Get the fundamentals right

So, we have a project. We are going to have developers contribute to it from a number of different, geographically scattered locations. How do you manage that? It was quite some months into the project before any serious code was written. A lot of time was spent planning what it was we were going to produce. What would it do? How would all the parts fit together? What was essential and what was optional? Would we adhere to standards - if so, which ones? There were many round-table discussions, many carried on into evening dinners and conversations in bars. What emerged was a fairly full specification for what a "Hydra object" looked like and, importantly, why. This fairly careful specification eventually meant that each partner had a good understanding of the joint goals and could work independently or with others towards achieving them - without too many crossed wires. In retrospect, some of those initial guidelines have been tossed aside as inappropriate to ongoing development but many have survived and still give commonality between many Hydra heads on both sides of the Atlantic.

The lesson for the Hydra Way: when you send people away from a meeting, make sure there is a clear understanding of what action will follow, and when. With widely scattered contributors, clarity is key. At the occasional more recent meeting we have sometimes forgotten that and felt the consequences.

Hydrangea

One of Hydra's less successful ventures (we're being honest here!) was Hydrangea. Once people got around to writing code we felt we needed a demonstrator and Hydrangea was it.

On the positive side, Hydrangea did briefly demonstrate what the basics of Hydra might deliver. It was simple and, interface-wise, rather basic - but it sort of worked.

On the negative side, it rapidly didn't work. Changes and improvements in the underlying code quickly broke it and Hydra's various developers were interested in pushing forward, well, development rather than maintaining a demonstrator.

It is perhaps regrettable that we didn't take it down sooner - but for the record, the lesson is not to commit to providing a non-core service that no-one has an incentive to maintain. For one thing, it looks bad when people find your broken offering. If you're going to do things in public view, Keep them current - and working!

Opening up the project

As time went on "real", production Hydra heads began to appear and others beyond the initial partnership became seriously interested in what we were doing. No longer was Hydra "vaporware", something presented at a conference as "in development", but here were real, live, accessible examples of Hydra in use. A few adventurous institutions started making noises about joining the party and we Hydra folks started to think that maybe we'd better get serious about formalizing such a process and making it attractive.

The idea of Hydra Partners was born: institutions who would commit to contributing to, and furthering, the community and its aims. Recognising that there were potential legal implications to expansion the partners jointly put in place a formal Memorandum of Understanding setting out such things as the aims of the project, its governance structure, and the terms of its intellectual property licensing and ownership. With this set out clearly, it was possible to invite additional Partners to join the Hydra project by signing a legal agreement to the MoU terms and providing a letter setting out what they envisaged their future contribution might be.

The process of writing the MoU and getting three university and two company counsels to agree it was lengthy. The process of getting further institutional counsels to ratify the Partner Agreement has proved to be likewise. That said, we are clear that this legal foundation was a Good Thing that now gives us surety that we and all our Partners understand the IP and licensing agreements around what we do. We later supplemented this Partner process by obtaining Contributor License Agreements from institutions and individuals to further ensure that the legal position is unambiguous.

Partners, Partners and Partners

To misquote, "if you build it, they will come." And come they did. Three or so years after opening up the project we are anticipating close on 30 partners and other production users at the end of the academic year. As we noted earlier, Partner meetings are held three or four times each year and all the partners are invited. As the group has got bigger these meetings have become more complex to orchestrate and we are currently looking hard at how best we should manage them as part of a wider strategy for holding together a community spread across the Atlantic and, probably soon across the globe.

The Partner meeting

So, how have we gone about holding Partner meetings? What has worked well? What continues to work well? What might we reconsider?

If we brand something as a Partner meeting then all Partners, new and old, are invited. We try to nail down the date of the next meeting before the end of the one preceding it so that people can take advantage of advance air fares; sometimes we manage this, sometimes we don't. At the very least we try to nail down the location and a timeframe. The invitation goes out to our Partner email list as soon as possible so that everyone knows about it. (More about our email lists later.)

All Partners, old and new, can expect a warm welcome and an inclusive meeting. As the meetings have got bigger we have found the need to have a facilitator for each round-table meeting session but, other than that, all people have an equal voice. Any newbies potentially have a valuable contribution to make as a fresh set of eyes. We've found it useful to "timebox" discussions so that they don't extend to take up more time than they really need. Generally the facilitator will agree an initial time block with the group (10/20/30 minutes) and at the end of that time the group will decide whether they're done with the topic or want to extend it for a further 10 minutes.

Generally, meetings nowadays have a twin track: sessions for "manager" types, and sessions for developers. Sometimes there are points of contact between the two tracks during the day, but sometimes not. Almost always the evening food and drinks is a good mix of all the folks in attendance. These evening sessions are an essential part of the meeting. We hesitate to use the term "bonding" - but you know what we mean.

The "managers" track tends to consist of a number of plenary sessions interspersed with breakout working groups. Hydra's strategic planning (in the widest sense) is generally handled in these sessions. Although the original Hydra Partners form a formal Steering Group, this is largely for administrative and legal reasons and it is the Partner meetings that shape Hydra's future. It used to be the case that we solicited agenda items on the wiki before a meeting and then pulled them into shape during the first plenary. That has worked less well as the meetings have got bigger and, perhaps reluctantly, we are now needing to do some structuring of the agenda in advance. It is clear that not all Partners can attend all meetings and so at the least a skeleton agenda is needed before travel arrangements are made: who should be sent, should we attend this meeting or save up to attend the next one? The finer detail of the agenda can still be worked out at the start of the meeting.

In the managers' track it is particularly important that someone be appointed "note taker" at the beginning of each session and that before it ends a set of action points, with names and dates against them, is established. Where we've forgotten to do this, or we've not done it too well, we've usually regretted it. The notes get published to the wiki as our formal record - part of our emerging philosophy that "if it isn't written down, it didn't happen".

The "developers" track has been arranged in a number of different formats: whatever it may be, it is taken as an opportunity for joint code planning and development. We are also now making an effort to provide an informal induction process for new developers who don't know the "old stagers" and who may need to get their heads around Hydra's way of doing things.

Communication

Between Partner meetings it is essential to the community that we stay in touch with each other so that Hydra initiatives can move along and Partners can keep track of what colleagues at other institutions are doing.

The Partners convene one group call each month (on the second Friday); in addition, the Steering Group also have a monthly call (on the fourth Friday). These calls are held on a conference line, generally at 08:30 Pacific, 11:30 Eastern, 16:30 UK, 17:30 central Europe. Hydra's developers generally have a call each Monday.

On a day-to-day basis, Hydra maintains a number of email lists:

Hydra-steering is a closed email list, largely for admin purposes.

Hydra-partners is a closed but general communication channel for Hydra Partners. The list is closed because such things as security vulnerabilities may be discussed there.

Hydra-tech is a general communication channel, mainly for the developers (who also tend to have an active IRC presence) (a user's first post is moderated to discourage spamming)

Hydra-announces is a general purpose, openly readable list, for general news announcements - gem releases and the like.

Web presence

Hydra maintains a web presence in two places. The "showcase" website is at projecthydra.org. This site attempts to give a new or potential user all the information they need to make a decision about investigating Hydra. It is fairly lightweight in content and carries information about the project generally, the software itself, and the community of Partners. We attempt to keep the newsfeed on the site reasonably busy so that it doesn't appear to become moribund.

More detailed information is held on the Hydra wiki at [Samvera](#) and many of the website's pages point into the wiki where things can be covered in more depth.

The Hydra software

Hydra's software is contributed by a considerable number of developers scattered across the globe. Hydra's dev community strives to be welcoming and inclusive. All the code is open source and available from github. New institutions are encouraged to try other institution's solutions before setting off on their own journey using the collection of Hydra gems. Anyone can submit proposed new code or patches: an inner core of developers, Hydra's "committers", will vet the code (perhaps suggesting improvements) before adding it to the appropriate gem.

One of Hydra's great strengths is that all our code is "test-driven". It contains a series of inbuilt tests that are run after any change to make sure that the change doesn't have unintended consequences and "break" something. We have a very high level of test coverage.