

DevMtg 2019-02-27

Developers Meeting on Weds, February 27, 2019

Today's Meeting Times



- DSpace Developers Meeting / Backlog Hour: 20:00 UTC in [#duraspace IRC](#) or [#dev-mtg Slack channel](#) (these two channels sync all conversations)

Agenda

Quick Reminders

Friendly reminders of upcoming meetings, discussions etc

- [DSpace 7 Working Group \(2016-2023\)](#): Next meeting is Thurs, Feb 14 at 15:00 UTC. Agenda: [2019-02-14 DSpace 7 Working Group Meeting](#)
- [DSpace 7 Entities Working Group \(2018-19\)](#): Next meeting on Tues, Feb 19 at 16:00 UTC
 - Last meeting notes at [2019-02-05 DSpace 7 Entities WG Meeting](#)
- [DSpace Developer Show and Tell Meetings](#): On hold until interesting topics arise.

Discussion Topics

If you have a topic you'd like to have added to the agenda, please just add it.

1. (Ongoing Topic) [DSpace 7 Status Updates](#) for this week (from [DSpace 7 Working Group \(2016-2023\)](#))
2. (Ongoing Topic) [DSpace 6.x Status Updates](#) for this week
 - a. 6.4 will surely happen at some point, but no definitive plan or schedule at this time. Please continue to help move forward / merge PRs into the dspace-6.x branch, and we can continue to monitor when a 6.4 release makes sense.
3. Meeting on Weds, March 3 (Next Week)
 - a. Another unfortunate conflict. Tim will be unable to attend (conflict with DSpace Leadership Meeting)
 - b. (This is the last such conflict in the Spring)
4. [Upgrading Solr Server for DSpace \(Mark H. Wood\)](#)
 - a. PR <https://github.com/DSpace/DSpace/pull/2058>
 - b. Docker configuration for external Solr
 - i. <https://github.com/Georgetown-University-Libraries/DSpace/commit/7115173d61776dd2455690518f5c9809cd0f28d4>
 1. The Dockerfile creates a new solr instance with 4 cores. It then overlays the schema and config changes in PR 2058.
 2. I attempted to create my branch so that I could create a PR back to Mark's branch, but some other changes from master seem to be showing up if I create a PR.
 - ii. This will need a small change to our docker compose files to invoke the external solr service. <https://github.com/DSpace-Labs/DSpace-Docker-Images/pull/79>
5. [DSpace Backend as One Webapp \(Tim Donohue\)](#)
 - a. PR: <https://github.com/DSpace/DSpace/pull/2265> (PR is in a reviewable state. SWORDv1 and SWORDv2 are merged into "Spring REST" webapp, with basic Integration Tests to prove both work)
6. [DSpace Docker and Cloud Deployment Goals \(old\) \(Terrence W Brady\)](#)
 - a. Build optimization PR reviews:
 - i. <https://github.com/DSpace/DSpace/pull/2344>
 - ii. <https://github.com/DSpace/DSpace/pull/2345>
 - iii. <https://github.com/DSpace/DSpace/pull/2346>
 - b. Add Docker build/push to Travis
 - i. This make sense to consider after 2307 is merged
 - ii. <https://github.com/DSpace/DSpace/pull/2308>
7. Brainstorms / ideas (Any quick updates to report?)
 - a. (On Hold, pending Steering/Leadership approval) Follow-up on "DSpace Top GitHub Contributors" site (Tim Donohue): <https://tdonohue.github.io/top-contributors/>
 - b. Bulk Operations Support Enhancements (from Mark H. Wood)
 - c. Curation System Needs (from Mark H. Wood)
 - i. PR 2180 improves reporting. Ready for review.
8. Tickets, Pull Requests or Email threads/discussions requiring more attention? (Please feel free to add any you wish to discuss under this topic)
 - a. Quick Win PRs: <https://github.com/DSpace/DSpace/pulls?q=is%3Aopen+review%3Aapproved+label%3A%22quick+win%22>

Tabled Topics

These topics are ones we've touched on in the past and likely need to revisit (with other interested parties). If a topic below is of interest to you, say something and we'll promote it to an agenda topic!

1. Management of database connections for DSpace going forward (7.0 and beyond). What behavior is ideal? Also see notes at [DSpace Database Access](#)

key	summary	type	created	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.								

5. Tickets requiring review. This is the JIRA Backlog of "Received" tickets:

key	summary	type	created	updated	assignee	reporter	priority
Unable to locate Jira server for this macro. It may be due to Application Link configuration.							

Meeting Notes

Meeting Transcript

Log from #dev-mtg Slack (All times are CST)

Tim Donohue [2:00 PM]

@here: It's our weekly DSpace DevMtg time. Agenda is at <https://wiki.duraspace.org/display/DSPACE/DevMtg+2019-02-27>

As usual, let's do a quick roll call to see who is able to join the meeting today

Mark Wood [2:00 PM]

Hi

Tim Donohue [2:01 PM]

Hi @mwood. Well, it might be a very quick meeting if it's just the two of us. :slightly_smiling_face: I'll go ahead and post some general updates here (just for anyone listening or looking in on notes later). After that, we can see if anyone else has popped in

First off, a general status update on DSpace 7 (for anyone interested). We are aiming for a 7.0 Preview Release (not all features, but many main ones) in late March. The Beta will follow that sometime in May (and Beta will be the first release including all DSpace features)

James Creel [2:04 PM]

Hey gang

Mark Wood [2:04 PM]

Hey, welcome.

Tim Donohue [2:05 PM]

These 7.0 timelines have slipped slightly cause of our volunteer workforce (we are all volunteers or on donated time, and sometimes that means features take a bit longer to create or review). However, we're working hard to ensure Beta is ready for user testing / training at OR2019.

Expect 7.0 Final sometime in July/Aug timeframe. The final release date will be easier to nail down after we run a community testathon (to see what bugs/issues still exist) on the Beta.

I think that's the 7.0 updates at a high level. Obviously though, for the nitty gritty details, check the DSpace 7 Working Group meeting notes or join a meeting.

Kim Shepherd [2:07 PM]

hi all, i'm semi-here

Tim Donohue [2:07 PM]

Hi James & Kim, welcome

Any questions on 7.0 updates?

Mark Wood [2:09 PM]

Bringing a beta to OR could be advantageous: a room full of testers with developers right there handy....

Kim Shepherd [2:09 PM]

sounds good to me

Tim Donohue [2:09 PM]

Oh, I should also mention there's a new Marketing Working Group just getting started. I know this is a DevMtg, but if you know folks interested in marketing/promotion, pass this along: <https://wiki.duraspace.org/display/DSPACE/DSpace+7+Marketing+Working+Group>

Ok, moving right along then.

Regarding 6.x updates (i.e. an eventual 6.4), the status hasn't changed much recently. I'll admit, all my effort is currently going towards helping get 7.0 out the door as soon as we can. I'm sure 6.4 will happen, but it's mostly waiting on someone (a Committer) to pick it up and run with it.

If anyone has more to add on a 6.4, or has any comments/questions, feel free to share them. That's the extent of the update though

James Creel [2:12 PM]

Not sure if this is exactly on topic, but I've got a librarian with a concern about support for deleted objects with OAI and it led me here: <https://wiki.duraspace.org/pages/viewpage.action?pageId=68064778>

and to this closed card: <https://jira.duraspace.org/browse/DS-2491>

Anyway, they would like "persistent" support for the deleted items which would mean some permanent tombstones for deleted stuff. Sounds like a fairly big architectural commitment. If we got a card for this on a sprint, might it go in to 6.x or 7's OAI?

If this seems like a bad idea I could communicate that

Kim Shepherd [2:14 PM]

a change in how OAI works wouldnt make it into a 6.x as it's not a bugfix

but could conceivably be part of 7.0 if we could figure out a good way to support it.. at the moment, expunged items are hard to keep track of on account of they really are completely gone :wink:

Tim Donohue [2:15 PM]

@jcreel256: that sounds like a significant change. Obviously, the reason OAI works the way it does is that the underlying object is gone when something is deleted. The only way to generate a "tombstone" is a withdrawal at this time.

Kim Shepherd [2:16 PM]

i've got some OAI improvements/fixes (both PMH and harvest) to contribute sometime, but it'll be a mix of fixes for 6.x and suggestions for 7.x (i've been working with harvesting a fair bit lately)

i'd suggest the librarian use withdrawing as the default 'deletion' method for now?

James Creel [2:16 PM]

Yes, makes perfect sense.

Tim Donohue [2:17 PM]

To be completely honest, the "window of opportunity" for 7.0 changes is closing rapidly. We are forced to really tighten our scope to ensure we hit that Beta goal for OR2019. So, anytime I hear of suggestions that sound major, I will be poking at them / pushing back on them.

Minor changes are still obviously welcome though

James Creel [2:17 PM]

Also sensible.

Mark Wood [2:18 PM]

Yes, I doubt there would be time available to handle a new 7.0 feature unless it is really simple, easy to review, and already written.

James Creel [2:19 PM]

On the bright side, the understanding that new features don't really have much of an opening right now could motivate those who want new stuff to help get the 7.0 out the door faster.

I can express this to my stakeholders around here.

Tim Donohue [2:20 PM]

Yes, exactly. :slightly_smiling_face: To be clear, 7.0 already has a ton of amazing new features. But, obviously, no release can do everything, and we're already actively limiting any additional "new features" from going into 7.0 (unless they are very small / self contained)

Mark Wood [2:21 PM]

And I've already started posting ideas for things to do in 8.0.

Tim Donohue [2:22 PM]

Thanks for that @mwood. I have been following your thinking/emails on dspace-devel. I just haven't found time to comment...but the ideas/brainstoms are all good

Mark Wood [2:22 PM]

Anyway, it sounds like this is something that DSpace could/should address.

Tim Donohue [2:22 PM]

agreed

So, this sounds like a wrapped up topic. Any other final notes/questions on 6.x?

Moving along then.

Topic #3: So, I noticed that (yet again) I have a meeting conflict next week. February's 28 days means this conflict came up both in Feb and now in March...but it won't happen again for the foreseeable future (I checked) So, this means, I won't be able to attend this meeting next week (March 3 at 15UTC). Anyone want to take chair responsibility next week? Or should we just cancel and touch base again on March 13? (edited)
Sorry, March 6
My days are totally wrong. I'll miss March 6 at 15UTC

Mark Wood [2:25 PM]
I can take the chair, if we think there is anything to discuss next week.

Tim Donohue [2:27 PM]
To be honest, the topics in this meeting have mostly been updates as of late. I'm not sure if that means this meeting is becoming *less useful* (in light of many other DSpace meetings in a week, mostly around DSpace 7), or if folks find it *more useful* (as it's one meeting to attend for a summary of what's going on).

Mark Wood [2:28 PM]
There are things (like the tombstone issue) that don't really fit into the more specialized meetings. And there are a number of WGs that will dissolve over the next few months so, fewer meetings....

Tim Donohue [2:28 PM]
But, beyond those usual "updates", I don't have specific topics for next week (just like I don't really have other specific topics this week). I do welcome anyone here to bring topics to the agenda though.

James Creel [2:29 PM]
Depends on your perspective, I'm sure. This meeting is sort of my speed since the 7 meetings would be on specific cards.

Kim Shepherd [2:29 PM]
if someone does pick up coordination of 6.4, this meeting is a pretty good place for 6.x release updates

Tim Donohue [2:31 PM]
Yes, to be clear, I'm not trying/impling canceling this meeting altogether. I'm just pointing out that this meeting is sorta a "weekly update meeting" of sorts right now. So, when it comes to specific topics for next week, I don't have any (other than I can pass along general updates to @mwood to report).
And I'm also wanting to note this "weekly update meeting" status out loud, as I've been unable to find/bring more specific topics to this meeting. But, I welcome others doing so in future weeks :slightly_smiling_face:

Mark Wood [2:33 PM]
It may be time to start thinking seriously about "what do we do after 7 releases?" and this would be a place to do it, if rapid interaction is needed.

Tim Donohue [2:33 PM]
And as @kshepherd notes, this definitely is the perfect meeting for 6.x questions/discussions. I've also tried to bring "backend" updates, like Solr upgrade & "One Webapp" refactor here...simply cause those are easy to also understand from a 6.x perspective.
@mwood: yes, that's true too. I haven't had much headspace for that quite yet :wink: But, I agree that the time is approaching
In any case, this was all a bit of a tangent.... The main question is whether to meet next week. @mwood if you want to lead, I can do my best to pass along topics (though ping me if I forget). If there's no specific topics /updates, you all could also just have a "bring your own topic" style meeting

Mark Wood [2:36 PM]
Do we want a meeting next week, then? If so, I can lead.

Tim Donohue [2:37 PM]
Since this is a "weekly update meeting" and I won't be there, I think it's a question of who will attend & what would you like to talk about. Or, just take the week off, if you want the hour to work on DSpace stuff :wink:
Honestly, either way is fine by me

James Creel [2:37 PM]
I can report back with my librarians' comments on the "transient" status of OAI responses for deleted items versus withdrawn ones
For what it's worth, the use case is our being harvested by EBSCO for EDS (EBSCO Discovery Service) and how they have some items in there that we deleted. The current ticket with them is still up in the air.
And who knows if they would even care if we had better compliance with OAI in this regard.

Tim Donohue [2:39 PM]
That seems like a fine topic to touch back on in more detail, if there's interest. Perhaps the goal could be to turn that use case into a JIRA ticket (that we can then work to schedule out)

Kim Shepherd [2:39 PM]

i'll be asleep at 15UTC so i'm neutral :wink:

Tim Donohue [2:40 PM]

Typical @kshepherd, always needing to sleep at 3am :laughing:
(Definitely don't get up and work at 3am. I never would)

Kim Shepherd [2:41 PM]

heh

Tim Donohue [2:41 PM]

Sounds like the meeting is "on", @mwood will lead it. You have a topic. Feel free to wrap it up early though if you run out of things to discuss. (And I'll catch up later in the logs)

Kim Shepherd [2:41 PM]

quick note from me regarding solr and docker change testing - i still haven't finished doing this properly sorry, but i'll update PRs when i can

Tim Donohue [2:42 PM]

Moving along now...yes, we can move into Solr upgrade updates from @mwood: <https://github.com/DSpace/DSpace/pull/2058>

Anything to update us on this week?

Mark Wood [2:42 PM]

No, I've been waiting to hear what folks think of it.

Tim Donohue [2:43 PM]

Ok, I'll admit, I haven't had time to look back at it myself either. I know Terry has been out a bit the last week or so. And Kim just said he hasn't had a chance.

Mark Wood [2:43 PM]

Understood.

I should go ahead and start tinkering with how to address existing sites, especially sharding.

Tim Donohue [2:44 PM]

So, it sounds like we're waiting on reviewers. I do also know that (in the DSpace 7 meeting last week) Art also volunteered to look at it. So, maybe we'll get an update on that tomorrow.

So, not hearing anything to discuss further here. I guess we'll move along

On the "One Webapp" backend side of things..<https://github.com/DSpace/DSpace/pull/2265> I'll note that as of *today*, I have OAI-PMH also working in the merged webapp. That means both SWORDS (v1 and v2) and OAI-PMH are all running alongside RESTv7 in Spring Boot.

The final piece of the puzzle is RDF...and a little bit of cleanup/tidying of the code

OAI-PMH proved to be the hardest piece so far, as OAI-PMH has its own (basic) integration tests. I had to move those over & convert them into Spring Boot compatibility. But, they all now work. And both SWORDS have basic ITs that also work

That's the update. The code is all there to have early reviews. You can even try it out if you wish. Feedback is more than welcome. I'm hoping RDF will prove to be a bit easier, and maybe this will be ready for official reviews in the next week.

Any questions or comments?

Mark Wood [2:50 PM]

Am I correct in thinking that, if one *did* still want to separate services (onto different boxes for instance) one could just run multiple copies of the Single Webapp with different bits enabled?

Tim Donohue [2:52 PM]

@mwood: kinda. However, currently you cannot turn "off" REST API v7. Everything else can be turned on/off though. So, if you deploy this webapp multiple times, you'd have multiple `/api` endpoints

Mark Wood [2:52 PM]

Ah. Thanks.

Tim Donohue [2:53 PM]

We might be able to look into treating REST API v7 similarly (and have the `/api` endpoint turn on/off), but I wasn't planning to do that as part of this initial PR.

Mark Wood [2:53 PM]

Sensible.

Kim Shepherd [2:55 PM]

Before we end, i just wanted to note that I might start a JIRA ticket or wiki page or something to discuss potential cleanup of xoai code and dspace-oai... i've run into a few things which i think could be refactored, eg. more standard/featureful xml I/O, better hibernate usage, some improvements to the base xoai xml format and

so on

so that could maybe be a place to discuss ways to support other OAI-PMH functionality, etc.

Mark Wood [2:55 PM]

One thing we ought to do at some point is move to XOAI v4.

Tim Donohue [2:56 PM]

@kshepherd: yes, that'd be welcome. We might want to start with a wiki page. I also noticed there's a bit of "crud" in the OAI-PMH codebase as I worked to pull it into Spring Boot (in that "one webapp" PR)
And yes, we really should look at moving to XOAI v4 (latest version) instead of still using XOAI v3.
(I think there's a ticket about moving to XOAI v4 already)

Kim Shepherd [2:57 PM]

@mwood yeh, that'd be part of it -- quite a lot of the stuff i want to fix is in xoai codebase rather than dspace-oai

yep <https://jira.duraspace.org/browse/DS-2595> is the "let's do xoai v4"

Tim Donohue [2:58 PM]

We are getting near the top of the hour, so I won't call any more topics (plus, since Terry is out today, it's hard to do Docker updates)

@kshepherd: yes, I'd like to get us up to XOAI v4 first (ideally) and then fix that version of the codebase. The way v3 works is a bit "odd" to me in some parts, and at least (at a glance) v4 looks cleaner (but I don't have experience to say whether that's actually true)

@kshepherd: in any case, a wiki page to *start* the discussion (and start gathering info/details/ideas) seems like a great place to begin.

Mark Wood [3:00 PM]

Quick question: DS-3989 is at +2 and has been for some time. There was talk about beefing up the tests, but it didn't block approval. Doing ITs for this kind of code is extremely difficult, and I'm inclined to merge this now and improve testing later. Thoughts?

Tim Donohue [3:01 PM]

Sorry, had to find the actual PR. This looks to be it: <https://github.com/DSpace/DSpace/pull/2180>

Kim Shepherd [3:01 PM]

sounds reasonable to me..

Tim Donohue [3:02 PM]

I'm OK with that direction, I see that @kshepherd had been assigned as a reviewer on this too...but Terry & Ben already gave it +1's. So, if everyone else is OK with it, I don't see a reason to hold it up.

Mark Wood [3:03 PM]

Thanks.

Tim Donohue [3:03 PM]

Although, it probably does need some documentation if it changes any user behavior. If it's all "underneath" refactoring, then that may not be necessary

Mark Wood [3:03 PM]

I will see to documenting it.

The Jira won't be closed until that's done.

Tim Donohue [3:04 PM]

Thanks! yes, that makes sense to move the Documentation part to JIRA and just note there that the ticket is waiting on final Docs to be created

OK, well, we are over time now (by ~5 mins). So, let's close up today's meeting.

Kim Shepherd [3:05 PM]

cheers all

Tim Donohue [3:05 PM]

I'll see you all (in this meeting) in two weeks! But, next week, @mwood will lead the meeting, you'll get an update from @jcreel256 on OAI-PMH persistence, and anyone else can bring topics too!
thanks all!

Mark Wood [3:06 PM]

'bye, all.

James Creel [3:06 PM]

Adios!

