

DevMtg 2018-10-31

Developers Meeting on Weds, October 31, 2018

Today's Meeting Times



- DSpace Developers Meeting / Backlog Hour: 15:00 UTC in [#duraspace IRC](#) or [#dev-mtg Slack channel](#) (these two channels sync all conversations)
 - Please note that all meetings are [publicly logged](#)

Agenda

Quick Reminders

Friendly reminders of upcoming meetings, discussions etc

- [DSpace 7 Working Group \(2016-2023\)](#): Next meeting was just before this one (no meeting tomorrow). Resuming on Thursday, Nov 8.
- [DSpace 7 Entities Working Group \(2018-19\)](#): Next meeting is Tues, Nov 6
 - Last meeting (including video recording): [2018-10-23 DSpace 7 Entities WG Meeting](#)
- [DSpace Developer Show and Tell Meetings](#): Met yesterday (Oct 30) on [Docker Updates](#)

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. DSpace 7.x dependency update PRs
 - a. <https://github.com/DSpace/DSpace/pull/2244> (Apache Commons Configuration v1. v2 and Commons Lang v3)
 - b. <https://github.com/DSpace/DSpace/pull/2241> (Log4j 1.2. 2.x) – Is this ready for review?
 - c. <https://github.com/DSpace/DSpace/pull/2058> (Solr v7).
4. DSpace 6x: Merge Legacy and Current Statistics: <https://github.com/DSpace/DSpace/pull/1810>
5. Docker PRs ([Terrence W Brady](#)) : [DSpace and Docker](#)
6. 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. Follow-up on Curation Task Reporting ([PR 2180](#))
 - c. [Bulk Operations Support Enhancements](#) (from [Mark H. Wood](#))
 - d. [Curation System Needs](#) (from [Terrence W Brady](#))
 - i. [PR 2181](#) implements per-run task parameters. Ready for review.
 - ii. [PR 2180](#) improves reporting. Ready for review.
7. 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](#)
 - a. In DSpace 5, each "Context" established a new DB connection. Context then committed or aborted the connection after it was done (based on results of that request). Context could also be shared between methods if a single transaction needed to perform actions across multiple methods.
 - b. In DSpace 6, Hibernate manages the DB connection pool. Each **thread** grabs a Connection from the pool. This means two Context objects could use the same Connection (if they are in the same thread). In other words, code can no longer assume each new `Context()` is treated as a new database transaction.
 - i. Should we be making use of `SessionFactory.openSession()` for READ-ONLY Contexts (or any change of Context state) to ensure we are creating a new Connection (and not simply modifying the state of an existing one)? Currently we always use `SessionFactory.getCurrentSession()` in `HibernateDBConnection`, which doesn't guarantee a new connection: https://github.com/DSpace/DSpace/blob/dspace-6_x/dspace-api/src/main/java/org/dspace/core/HibernateDBConnection.java

Ticket Summaries

1. Help us test / code review! These are tickets needing code review/testing and flagged for a future release (ordered by release & priority)

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

2. Newly created tickets this week:

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

3. Old, unresolved tickets with activity this week:

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

4. Tickets resolved this week:

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 CDT)

Tim Donohue [10:00 AM]

@here: it's DevMtg time. Agenda at: <https://wiki.duraspace.org/display/DSPACE/DevMtg+2018-10-31>
Let's do a quick roll call to see who is able to join today

Terry Brady [10:00 AM]

hello

Mark Wood [10:00 AM]

Hi

Alexander Sulfrian [10:02 AM]

Hi

Tim Donohue [10:03 AM]

Hi all, welcome. So, I admit, I don't have a ton to update us on in the "Quick Reminders" today. I think the notes there pretty much speak for themselves.

Main points are the next Entities WG (#entities-wg) meeting is next Tues, Nov 6 (agenda coming soon). Also, we had a Dev Show & Tell on Docker yesterday. Video at <https://wiki.duraspace.org/display/DSPACE/Dev+Show+and+Tell+--+Oct+30%2C+2018+--+1500UTC+--+DSPACE+Docker+Updates>

Terry Brady [10:04 AM]

Note, there are no additional Show and Tell meetings scheduled. We can schedule them as the interest arises. It was a good meeting yesterday, but I was surprised that the meeting did not draw a larger crowd.

Tim Donohue [10:04 AM]

I'm also going to skip over the "ongoing topics" on DSpace 7 and 6....as we have other agenda items that talk to both of these...so, let's skip to topic #3.

@terrywbrady: yes, it was disappointing that only ~5 folks were there, but hopefully others will catch the video. I still think it was a nice update on the Docker work

On to topic #3, I'm keeping a running list of dependency update PRs in the works for DSpace 7

As noted last week, I'd like to see this group help us move some of that work along, as most of this effort is simply updating Java dependencies (which all of you should be familiar with, as we've had these same dependencies in DSpace 6 and below)

So, of note specifically...

<https://github.com/DSpace/DSpace/pull/2244> (Commons Config upgrade)

DSpaceSlackBot (IRC) APP [10:07 AM]

AlexS[zemat] has quit the IRC channel

Tim Donohue [10:08 AM]

This PR could use additional reviews. It has very extensive Unit Tests to prove the upgrade works. But, I need another +1

Mark Wood [10:08 AM]

Needs one more positive review.

Tim Donohue [10:08 AM]

The next two PRs are still a "work in progress" (I believe), but worth paying attention to & perhaps giving early feedback on

<https://github.com/DSpace/DSpace/pull/2241> (Upgrade of log4j)

<https://github.com/DSpace/DSpace/pull/2058> (Upgrade of Solr to v7)

Mark Wood [10:09 AM]

I need to get logging working again in dspace-spring-rest and then 2241 will be ready for review. Comments are always welcome.

Tim Donohue [10:10 AM]

@mwood: do you need any help in making that happen? Or do you essentially have a direction to move 2241 forward?

Mark Wood [10:10 AM]

I have a direction to try. If I can't get it finished soon, I'll ask for help.

Tim Donohue [10:11 AM]

Sounds good. Definitely let me know if you get stuck...and I may be able to give a second set of eyes to any issues

Mark Wood [10:11 AM]

Spring Boot is configured differently than every other webapp. in the universe, so I've had some learning to do.

Tim Donohue [10:12 AM]

Yes, Spring Boot has a bit of "convention over configuration" to it.... so, there's some things that happen a bit "magically" until you learn how it all works. That said, I've grown to love Spring Boot myself (as I think it's easier than traditional webapps in many ways)...but you have to understand that it does do things a tad differently at times.

So, if you hit Spring Boot issues, I'm glad to help.

Mark Wood [10:13 AM]

Thanks!

Tim Donohue [10:14 AM]

As for the Solr upgrade.. I'll note I've added some commits to the Solr PR recently (and actually turned off some Spring Boot autoconfiguration "magic" with regards to Solr that caused us issues)

That PR is getting "close" but still has one remaining Unit test failure. I may have time to look closer at it in the next few days

Mark Wood [10:14 AM]

I saw that come in, and I need to spend some time on it.

Tim Donohue [10:16 AM]

As a part of the Solr upgrade though, I hit immediately on our Solr Statistics Schema issues (which is the next topic on the agenda). So, this might be a good segue to looking at that a bit more...

Specifically, this PR: <https://github.com/DSpace/DSpace/pull/1810> (From @terrywbrady)

Mark Wood [10:17 AM]

I have managed to snapshot a DSpace 5 statistics core for testing. Now I just need to do the testing.

Terry Brady [10:17 AM]

I have this next on my DSpace TODO list... to re-test and cleanup this code based on the review suggestions.

Tim Donohue [10:17 AM]

I think we have two topics with regards to PR#1810... (1) we need a tester/reviewer (or two) to help us verify the migration process. (2) We need to discuss whether to backport this to 6.x for a 6.4 release

Terry Brady [10:18 AM]

It is a tricky thing to test since you need to prep a DSpace 5 instance for upgrade to DSpace 6.

Tim Donohue [10:18 AM]

I'm able to offer code review help, but I don't have a decent data set to test with.

And, I'm of the opinion that this *should be backported to 6.x* (for a 6.4), once we get it tested/merged into `master`. It should be a simple cherry-pick though, as it's new code

Mark Wood [10:20 AM]

Sensible.

Tim Donohue [10:21 AM]

Is there anyone else @here who could help us test/verify the migration script in <https://github.com/DSpace/DSpace/pull/1810> ? This seems like a very useful fix for people annoyed by this in 6.x...we really could move it along quickly with feedback

If you are willing, just please let us know in the PR comments -- or test and report back in the PR comments. I think that's it on that topic.

The remaining topics on the agenda have been there for some time...and I'm not sure whether we have any real updates to share, but we can see

@terrywbrady: any updates on Docker you want to share? (topic #5)

Terry Brady [10:24 AM]

Nope. The PR's are merged and the changes have been very handy. Feel free to remove this from the agenda. I imagine we'll get some enhancement suggestions as folks try out the images.

Tim Donohue [10:25 AM]

Ok, I'll take that off the agenda

Next up, we have a variety of "Brainstorms / Ideas" under topic #6. I suspect many of these have been set aside temporarily (as other work has ramped up), but wanted to check in.

For my update, the "DSpace Top GitHub Contributors" site is still being maintained...but I haven't had a chance to get it on the Steering/Leadership group agenda...they are discussing more important things like DSpace 7 right now

(And I want to give them a chance to give feedback before it becomes anything "official". So, right now, it's just "unofficial", but feel free to check in on it)

Terry Brady [10:27 AM]

Tim, it would be great to approve your solution if you add it to <https://stackoverflow.com/questions/51972020/github-graphql-query-count-pr-reviews-by-user-for-a-given-month>
I suspect other folks would benefit from what you discovered.

Tim Donohue [10:28 AM]

I'll try to do that. Unfortunately, the answer is essentially "you cannot do that in GraphQL"...you need to use another tool

Terry Brady [10:29 AM]

OK... I guess I did not look at your solution close enough.

Tim Donohue [10:29 AM]

But, I can try and find time to summarize & link back to the top contributors list code
My solution was essentially to pull in the data of **all reviews** (in a massive list) from GraphQL, but then I had to parse, sort, group them into a "Top 10" list using another tool
Are there any other updates to that topic #6 set of Brainstorms?

Mark Wood [10:32 AM]

Curation PRs are still looking for additional review.

Terry Brady [10:33 AM]

I have approved them. These will be nice to build upon.

Mark Wood [10:33 AM]

Thank you for reviewing.

Tim Donohue [10:34 AM]

Ok, I'll see if I can find some time to review these too. I admit, I haven't looked at them in a while, but should do so

Mark Wood [10:35 AM]

Thank you.

Tim Donohue [10:35 AM]

I think that basically brings us to the end of our agenda. Are there other topics, brainstorms or PRs that anyone would like to bring up for discussion?

Terry Brady [10:36 AM]

For 6d, it might be good to list @mwood as the point of contact on the agenda rather than me.

Tim Donohue [10:36 AM]

Oh, sure, I think that never got updated

One thing of note...I've noticed a **lot** of Google Analytics PRs coming in from Atmire. So, if anyone is interested in helping enhance the integration with GA, you might want to look at these:

- <https://github.com/DSpace/DSpace/pull/2248>
- <https://github.com/DSpace/DSpace/pull/2251>
- <https://github.com/DSpace/DSpace/pull/2252>

All these PRs sound interesting to me (just at a glance), but I've not been able to review or test them yet. Again, I could use more eyes / reviewers / testers to help move these along

Terry Brady [10:40 AM]

We tried to integrate some of the GA stuff with our DSpace 5 instance a year ago and we hit a wall with it. Our current plan is to retry this after we upgrade to DSpace 6. I know there is a lot of interest in querying all of our stats in one place rather than relying on a separate DSpace stats report tool.

Tim Donohue [10:41 AM]

I'll also note, I've recently started brainstorming ways to simplify our DSpace 7 installation process. I admit, I don't have any amazing ideas or code to share yet...but I want to see if we can simplify what is becoming a more complex install process.

My initial direction was in <https://github.com/DSpace/DSpace/pull/2231> ... but that got some negative feedback from DSpace 7 team. So, I'm looking towards other ideas.

(I'm just saying this here to keep me accountable, and to let folks know you can throw ideas at me, if you have any,)

Mark Wood [10:42 AM]

I have a back-burner project to build an "installer" artifact.

DSpaceSlackBot (IRC) APP [10:43 AM]
renilgh has quit the IRC channel

Mark Wood [10:43 AM]

One of those wizard things that quizzes you about how you want the install done and then pulls the pieces from a local or remote Maven repo.

Tim Donohue [10:43 AM]

@mwood: interesting. Admittedly, that was the direction I was initially going towards with the "one webapp" backend -- as if we put everything into a single Spring Boot webapp, Spring Boot has it's own tools in this area.

But, it seems like others aren't interested in that direction, so I'm starting to dig around into other ways of achieving this...both for backend *and* for Angular frontend

Mark Wood [10:44 AM]

I have some things I want to do, like excluding the DBMS drivers because I'll provide DBMS pools out of JNDI. It still uses the Ant script to drive the process, but builds Ant and Ivy into the tool. No more need to *install* Ant.

Tim Donohue [10:46 AM]

That's another thing I've look towards... using Ant but not needing it be installed. The same thing for Maven, which technically could be replaced by `mvnw` (Maven Wrapper, a tool that installs Maven for you if you don't have it already: <https://github.com/takari/maven-wrapper>)

(Maven Wrapper is another thing I found from Spring Boot...as it's used by Spring Boot by default if you create a new project.)

Terry Brady [10:46 AM]

The ant installation was one of the most complex parts of creating the docker image for DSpace. It would be nice to have that handled within the code.

Mark Wood [10:48 AM]

Per Hardy Pottinger, the installation process should also produce an answers file so that you can re-run it noninteractively. It shouldn't be hard to do.

Tim Donohue [10:49 AM]

In any case, this is an area of interest to me... and if it interests others of you too, we could keep touching on this (in this meeting) and see if we can simplify this installation process together.

Mark Wood [10:49 AM]

One could open an editor and just write an answers file....

I do think that DSpace needs a better-orchestrated installation process to help out newer users.

Tim Donohue [10:50 AM]

@mwood: While that would be nice, I think there are smaller steps we could take that would make a big difference. My fear of a custom noninteractive installation is that it becomes a pain to support cross-OS, unless you find a third-party tool that already does that.

So, I'm specifically looking for best practices we can "borrow" from other projects (like Spring Boot) that can modernize our installation process.... and then hopefully we can improve further on that iteratively

Terry Brady [10:51 AM]

On a side note, it seems like we are seeing an uptick in tech support questions. I presume that signals lots of DSpace adoption and folks anticipating help from the community. I am finding it difficult to keep up and respond as I would like.

Mark Wood [10:51 AM]

Hm, I don't like the way Spring Boot tends to reinvent things I already know how to do....

Tim Donohue [10:53 AM]

@mwood: Spring Boot actually doesn't "reinvent" all that much overall -- it tends to borrow a lot of what it does from other Spring projects. It's more a collection of "best practices" to make building / installing Spring webapps easier.

And the reason I call it out specifically is that we are now moving in that direction of everything Spring... so, it might be nice to follow their "best practices" where we can do so

@terrywbrady: yes, I've been trying to keep up in those questions myself (to hopefully keep others concentrating on DSpace 7). But, I agree, there's a lot coming in these days

Mark Wood [10:54 AM]

would like to throw away the ServiceManager and let Spring manage services. Since we use it for that anyway.

Tim Donohue [10:55 AM]

:+1: to that idea. I admit, it just hasn't been a high priority yet for me (as it "works" even though it's not ideal)

But, I'd support moving everything into Spring and avoiding that layer of unnecessary abstraction

So, we are heading to the top of the hour. This has been good discussion (and I'll add this "easy install" idea to our agenda, and cleanup other topics a bit). Are there any last notes/comments to add before we close up?

Mark Wood [10:58 AM]

You mentioned @TODO in a PR comment. There isn't a standard such thing that I know of, but it brings up an idea: "global" TODOs that should be noted when one is not actually in the code so marked. IDEs handle "local" TODOs just fine, but not this.

Does that make sense?

I found this: <https://github.com/cowlingj/java-TODO-annotation> which might be made into e.g. a Maven reporting plugin.

Terry Brady [10:59 AM]

Have a good week

Tim Donohue [10:59 AM]

A "global" TODO is just a JIRA ticket, is it not?

Mark Wood [10:59 AM]

Just bringing it up because I have the ear of some fellow Java dev.s

@tdonohue maybe so.

Something to think about.

Tim Donohue [11:02 AM]

Yes, in any case, a JIRA ticket is probably the better overall "TODO". But, I sometimes duplicate these notes into code as the IDE does let you pull out a list of all TODOs (at least most IDEs do)

Mark Wood [11:02 AM]

I will poke around in NetBeans to see if I've missed something.

Tim Donohue [11:02 AM]

Ok, we'll wrap things up for today. Thanks all! Talk to you next week at 20UTC (and reminder that with daylight savings, this will be one hour *earlier*)

Mark Wood [11:02 AM]

Thanks, all!