

DevMtg 2010-12-01

Developers Meeting on Weds, December 1, 2010

Agenda

- [DSpace 1.7](#) Status Updates & Discussion
 - Any 1.7.0 JIRA issues requiring our immediate attention/discussion? (**Bring specific issues to meeting, or add to this agenda!**)
 - [DS-613](#) - would appreciate any thoughts on attached patch. Is it "good enough" for 1.7.0?
- Upcoming Release Schedule:
 - December 3, 2010 : Release Candidate 2
 - December 6-15, 2010 : Final Testing / Bug Fixing
 - December 17, 2010 : Final Release
- [JIRA Catch-Up](#) (time permitting) – starting with issue DS-641
 - [Search for recent JIRA Issues](#)

Meeting Summary

1.7.0 updates & schedule reminder, demo.dspace.org memory usage

- 1.7.0 Updates
 - DS-613 – will be reviewed by Robin Taylor
 - DS-750 – will be reviewed by Tim Donohue
 - RC2 will be cut on Friday or over the coming weekend.

Meeting Transcript

- Our normal IRC Logger wasn't working today. Full transcript follows:

```
(1:58:31 PM) tdonohue: DSpace Devel mtg to start here in a few minutes. Today's agenda: https://wiki.duraspace.org/display/DSPACE/DevMtg+2010-12-01
(1:59:04 PM) HardyPottinger [80ce8627@gateway/web/freenode/ip.128.206.134.39] entered the room.
(1:59:42 PM) cccharles [~chris@131.104.62.55] entered the room.
(2:01:04 PM) tdonohue: Hi all -- it's top of the hour, which means it's time for the DSpace Devel mtg. Today's agenda (for those who just joined up): https://wiki.duraspace.org/display/DSPACE/DevMtg+2010-12-01
(2:02:08 PM) tdonohue: Only real topic is 1.7.0 updates -- especially checking in with everyone who has outstanding JIRA issues, and whether or not there are any that need our immediate attention
(2:02:39 PM) richardrodgers [~richardro@dhcp-18-111-21-39.dyn.mit.edu] entered the room.
(2:03:25 PM) tdonohue: Currently, we still have 42 unresolved issues for 1.7.0. Any that are unassigned will likely be rescheduled for 1.8.0 (unless claimed very soon): https://jira.duraspace.org/secure/IssueNavigator.jspp?reset=true&jqlQuery=project+%3D+DS+AND+resolution+%3D+Unresolved+AND+fixVersion+%3D+%221.7.0%22+ORDER+BY+priority+DESC&mode=hide
(2:04:19 PM) tdonohue: PeterDietz -- do you know of any JIRA issues that need immediate review for 1.7.0? Does anyone else have any they'd like us to review?
(2:05:33 PM) ***tdonohue just realized PeterDietz isn't here
(2:06:28 PM) tdonohue: Well, I have one issue I'd like a "quick review" on. DS-613 (https://jira.duraspace.org/browse/DS-613) I have a very simple solution (see the patch), but it's not perfect.
(2:06:53 PM) tdonohue: thoughts/comments welcome
(2:07:27 PM) richardrodgers: +1 if it's been tested
(2:07:27 PM) richardrodgers left the room (quit: Read error: Connection reset by peer).
(2:07:55 PM) robint_: +1. I'll test it if you want further testing.
(2:08:23 PM) tdonohue: I'd appreciate further testing -- I've done my own local testing, and it seems to work fine, but would like a second opinion :)
(2:08:49 PM) robint_: I should have some time tomorrow.
(2:08:50 PM) keithg [~keith-noa@lib-kgilbertson.library.gatech.edu] entered the room.
(2:08:54 PM) tdonohue: so, robint_, I'd appreciate if you could test it and add comments. thanks!
(2:09:35 PM) tdonohue: Anyone else have any open 1.7.0 issues they'd like a quick review on, or need help on?
(2:09:46 PM) richardrodgers [~richardro@dhcp-18-111-21-39.dyn.mit.edu] entered the room.
(2:10:17 PM) jefftrimble [~jefftrimb@maag127.maag.ysu.edu] entered the room.
(2:10:42 PM) tdonohue: repeating for richardrodgers and jefftrimble, who just popped in: "Anyone else have any open 1.7.0 issues they'd like a quick review on, or need help on?"
(2:11:07 PM) jefftrimble: <==feels special.
(2:11:11 PM) tdonohue: :)
(2:11:20 PM) richardrodgers: for tdonohue: have you had a look at the curation admin UI fixes?
(2:12:02 PM) tdonohue: richardrodgers -- hadn't had a chance yet, actually. At a glance, looked good. But, I'll make sure to do some actual testing tomorrow or Friday
(2:12:17 PM) richardrodgers: OK - thanks
```

(2:13:15 PM) richardrodgers: question - I thought we were cutting an RC2
(2:14:08 PM) tdonohue: Yes -- as far as I'm aware, Peter is cutting an RC2 on Friday (Dec 3) -- that's the schedule that PeterDietz was following on the 1.7.0 release notes: <https://wiki.dspace.org/display/DSPACE/DSpace+Release+1.7.0+Notes>
(2:14:33 PM) tdonohue: (unfortunately, PeterDietz isn't around now to verify that -- but, last I heard from him, that was the plan)
(2:15:25 PM) richardrodgers: Cool - I just forgot the date. But that means we should try to indentify and fix priority bugs before then
(2:16:47 PM) tdonohue: Yes -- I'd definitely encourage everyone to try and get bug fixes squared away before Friday (or let PeterDietz & I know if something will be delayed much longer). As of next week, we'll also make sure that demo.dspace.org gets updated to running RC2 (so we can bang on it a bit more before final release)
(2:17:23 PM) tdonohue: Any other questions/comments on the upcoming 1.7.0 release process/schedule, etc.?
(2:19:23 PM) tdonohue: Ok -- guess not :)
(2:20:12 PM) tdonohue: Slightly related topic: As committers are already aware, demo.dspace.org has run into some various OutOfMemory exceptions this past week while running 1.7.0RC1
(2:20:44 PM) richardrodgers: Do we think we understand it yet?
(2:20:56 PM) tdonohue: An Update: we just hit another "OutOfMemory error: permgen space" exception about an hour ago. This is now 3 days in a row that we've seen that.
(2:21:21 PM) richardrodgers: I know you had some Tomcat memory config issues/revelations....
(2:21:57 PM) tdonohue: I got a snapshot from YourKit (which was automatically generated at the time of the OutOfMemory exception). If anyone has time or willingness to help me try and diagnose these issues, I'd appreciate help.
(2:22:34 PM) richardrodgers: I could take a crack at it. Are you configuring permgen specifically?
(2:22:42 PM) tdonohue: richardrodgers -- yea, I did have one revelation -- that Java Heap space was less than I thought -- we only had 256MB allocated instead of recommended 512MB. However, this didn't fix permgen problems
(2:22:55 PM) tdonohue: no -- permgen is just at it's default of 64MB
(2:23:12 PM) richardrodgers: Could definitely bump that
(2:23:27 PM) tdonohue: Ok -- I'll find a place to post up this YourKit Snapshot (it's large) for others to look at.
(2:24:29 PM) tdonohue: richardrodgers -- actually, that's part of the question here too. Obviously, I could bump it, but hasn't DSpace run fine before without bumping up Permgen? Otherwise, we should change our recommendations for installing DSpace in production to recommend bumping up PermGen
(2:24:29 PM) richardrodgers left the room (quit: Read error: Connection reset by peer).
(2:25:18 PM) HardyPottinger: Mark Wood sent out his settings a while back to dspace-tech, and we've implemented them, permgen is set to 128M
(2:25:46 PM) HardyPottinger: MaxPermSize, actually
(2:25:58 PM) tdonohue: HardyPottinger -- aha, that's good to know. Do you have a copy of his email?
(2:26:28 PM) richardrodgers [-richardro@dhcp-18-111-21-39.dyn.mit.edu] entered the room.
(2:26:40 PM) robint_: While I was watching the demo site today Permgen was close to its limit all today which suggests its like that from deployment.
(2:26:55 PM) HardyPottinger: tdonohue: I can track it down, pretty sure I put a copy in our wiki
(2:27:34 PM) richardrodgers: tdonohue: (I'm having connectivity probs, sorry if I repeat myself) Why not bump permgen to 128M?
(2:27:35 PM) tdonohue: thanks HardyPottinger -- just wanted to compare Mark's settings with what I have. We may need to provide better recommendations on this directly in the DSpace installation Documentation
(2:28:44 PM) tdonohue: richardrodgers: yea, we could do that -- I just wasn't sure if we want to then start *recommending* in Dspace Docs that everyone run Dspace with 128M permgen. I'm trying to run demo.dspace.org as sort of a "default install" to get a better sense of what our real recommendations are around memory, etc
(2:29:21 PM) richardrodgers: understood tdonohue , I just meant initially as a diagnostic....
(2:29:23 PM) tdonohue: So, as an example, we already discovered DSpace minimally needs about 512MB heap space memory allocated (as we hit some errors at only 256MB)
(2:30:01 PM) tdonohue: richardrodgers -- yea, that is true -- we could bump it up to see if it happens to keep climbing, or if it just happens to always hover around 60-64MB default value
(2:30:21 PM) tdonohue: So, I'll go ahead and bump it up today, and we can watch what happens
(2:30:25 PM) richardrodgers: yep, exactly
(2:33:06 PM) tdonohue: Ok -- I actually didn't have anything else specific to discuss today. Anyone have anything else they'd like to bring up?
(2:34:05 PM) richardrodgers: Guess we can't thrash over Git without Peter ;)
(2:36:04 PM) HardyPottinger: tdonohue: link to Mark's note on Tomcat configs: <http://www.mail-archive.com/dspace-tech@lists.sourceforge.net/msg12075.html>
(2:36:20 PM) tdonohue: richardrodgers: yea, an indepth Git discussion would probably be good to include Peter. We could still toss that around a bit now, but I'm noticing a lack of activity in this discussion already :) We only seem to have 3-4 people active here
(2:36:47 PM) tdonohue: Thanks for tracking that down HardyPottinger. I'll take a glance at what he is recommending
(2:37:20 PM) richardrodgers: I think you are right - by the way, I like your suggestion of learning from the Fedora experience...
(2:38:04 PM) HardyPottinger: our experience is the same, that PermGen fills just a bit over half and hovers there, so just a tad over 64
(2:39:04 PM) tdonohue: richardrodgers-- yea, I think the Git idea is an interesting one. But it would be good

to see how Fedora handles it and what they experience. I know already they are already debating whether to move entirely to Git, or whether to continue to "sync" back to a Subversion copy somewhere

(2:40:00 PM) tdonohue: HardyPottinger: also good info. That's actually similar to what is happening on demo.dspace.org right now. it hovers around 60MB though, and eventually bumps up against 64MB limit at which point it throws OutOfMemory errors. We'll have to see what happens once I bump it up to 128MB

(2:42:06 PM) HardyPottinger: tdonohue: we've also implemented Bill Anderson's (GATech) recommendation to use the parallel garbage collector, still tracking down that message

(2:42:07 PM) richardrodgers left the room (quit: Read error: Connection reset by peer).

(2:43:38 PM) tdonohue: Ok -- since we only seem to have 3 of us in here (everyone else is quiet!), do we just close up early? Is there anything else anyone needs/wants to discuss?

(2:43:54 PM) kshepherd: i'm here, i just have nothing to add ;)

(2:44:23 PM) richardrodgers [~richardro@dhcp-18-111-21-39.dyn.mit.edu] entered the room.

(2:44:51 PM) tdonohue: jefftrimble, is there anything to discuss around 1.7.0 Docs? (just wanted to check in, since you are here)

(2:45:07 PM) robint_: I'm lurking too but I've got to go now. Cheers all.

(2:45:27 PM) richardrodgers: bye Robin

(2:45:43 PM) kshepherd: except to note that at UoA we use a 2GB heap and 256MB permgen..

(2:45:44 PM) tdonohue: bye robint_ thanks for volunteering to look at DS-613

(2:45:47 PM) kshepherd: cheers robin

(2:46:38 PM) tdonohue: kshepherd -- ok, good to know.

(2:46:42 PM) HardyPottinger: tdonohue: here's the link to Bill's note that mentions the ParallelGC, but he doesn't really mention it in detail, it's just in the config line, I remember doing further reading and decided to implement it: <http://www.mail-archive.com/dspace-tech@lists.sourceforge.net/msg12086.html>

(2:47:44 PM) jefftrimble: No.....

(2:47:47 PM) jefftrimble: sorry...

(2:48:01 PM) jefftrimble: just cleaning up and working away.

(2:48:18 PM) jefftrimble: We're still learning about the Confluence software and its ins and outs....

(2:48:20 PM) tdonohue: jefftrimble -- nothing to be sorry about. Just trying to determine if you need help on anything, or could use some extra "proof readers" or anything.

(2:48:32 PM) jefftrimble: Proof readers are ALWAYS welcome.

(2:48:53 PM) jefftrimble: I can always tell if someone messes something up...and I don't have to beat anyone if they do. :-)

(2:49:30 PM) tdonohue: Ok -- if anyone can take some time to help us proof read the Documentation in coming weeks, it'd be much appreciated! I'll make sure to help out as well

(2:50:08 PM) tdonohue: (the more proof reading we can do now, the better!)

(2:50:38 PM) tdonohue: HardyPottinger: thanks for that other note from Bill. i'll reread that

(2:50:54 PM) grahamtriggs: I run a 4GB heap, and a large permgen (can't remember whether it's 512 or 1024mb off hand)... but that's got 30 repositories in it, and has lots of free space (about half of the heap)

(2:53:47 PM) richardrodgers_ [~richardro@dhcp-18-111-21-39.dyn.mit.edu] entered the room.

(2:53:47 PM) richardrodgers left the room (quit: Read error: Connection reset by peer).

(2:53:47 PM) richardrodgers_ is now known as richardrodgers

(2:55:16 PM) tdonohue: ok -- also good to know. Sounds like we have some relatively common settings amongst everyone, which could lead us towards some "best practice" or "recommended" configs for DSpace (at least in terms of memory allocation)

(2:57:13 PM) tdonohue: Well, since we're nearing the 1 hour mark, I think I'll just close up a bit early. I'll still be around here for a while though, if anything else comes up. I'll also go bump up demo.dspace.org to 128MB permgen very shortly to see what happens

(2:57:42 PM) richardrodgers: Thanks all

(2:57:46 PM) grahamtriggs: (btw - it's also been running for over 120 days without restarts/reboots. Individual repositories have been hot undeployed/redeployed a number of times to push changes live. Stability and resource usage is a choice :)

(2:58:08 PM) richardrodgers left the room (quit: Client Quit).

(2:59:17 PM) kshepherd: grahamtriggs: largely JSPUI?

(3:01:31 PM) grahamtriggs: all JSPUI... in my experience, you wouldn't get the xmlui resource usage remotely close... and you wouldn't be able to do the hot undeploy / redeploy (at least not repeatedly)... I can't really comment on long term stability

(3:01:34 PM) tdonohue: FYI -- I just bumped up demo.dspace.org to -XX:MaxPermSize=128M

(3:01:45 PM) kshepherd: grahamtriggs: thought so, just checking ;)

(3:08:58 PM) PeterDietz [~PeterDiet@ACK5859s3.lib.ohio-state.edu] entered the room.

(3:09:01 PM) PeterDietz: hi all

(3:13:17 PM) kshepherd: g'day PeterDietz

(3:14:03 PM) PeterDietz: Hi, did I completely miss the meeting?

(3:14:33 PM) kshepherd: kinda ;)

(3:15:21 PM) PeterDietz: great, there's no log online: <http://irclogs.duraspace.org/index.php?date=2010-12-01> .. would you mind pastbin'ing one to me?

(3:15:30 PM) kshepherd: was a very quiet/casual meeting, mostly discussing resource usage on demo.dspace.org and elsewhere, quick schedule rehash, quick review of DS-613

(3:15:33 PM) kshepherd: sure

(3:15:54 PM) kshepherd: actually, that's one thing i did want to ask opinions on

(3:16:11 PM) tdonohue: ack -- looks like our logger wasn't working today, even though it's supposedly in this room. What's up, DuraLogBot?

(3:16:50 PM) kshepherd: sometime i'd like to make a cool IRC bot that could interact with JIRA, SVN, Fisheye, (and maybe even the demo site via some easy interfaces;) etc for #dspace and #duraspace

(3:16:55 PM) kshepherd: what do people think?

(3:17:07 PM) kshepherd: so we could do @vote DS-613

(3:17:11 PM) PeterDietz: The "open source" club at my school/univ built a pretty awesome ircbot.. its fun to interact with

(3:17:28 PM) kshepherd: and the bot would display the description, last comment, and hold a vote for X minutes

(3:17:51 PM) kshepherd: PeterDietz: the supybot in #code4lib is pretty impressive, even if most of its cool features are just timewasters ;)

(3:18:10 PM) kshepherd: i'd particularly like to just see a stream of svn activity in #dspace

(3:18:15 PM) kshepherd: just an idea anyway

(3:18:44 PM) PeterDietz: or when someone say DS-666, the bot would change it to a link <http://jira.dspace.org/.../DS-666>

(3:20:23 PM) tdonohue: that would be neat -- I've also seen the #code4lib bot

(3:21:21 PM) kshepherd: PeterDietz: <http://dspace.pastebin.com/4Sqc9996>

(3:21:30 PM) kshepherd: not a perfect paste, sorry

(3:21:37 PM) robint_ left the room (quit: Quit: Page closed).

(3:22:30 PM) kshepherd: PeterDietz: yeah, just handy stuff like that for IRC addicts like me to get realtime "dspace stuff" and help with the voting and jira reviews and recent svn changes, make clickable links of everything, even search the mailing list archives

(3:22:52 PM) kshepherd: i can't figure out a nice way of doing svn stuff without a 'push' process of some kind running on scm.dspace.org, but i'll keep thinking

(3:23:24 PM) tdonohue: kshepherd: if you ever get around to building such an amazing IRC bot, I think DuraSpace would be willing to set it up to run in #duraspace and/or #dspace :)

(3:24:00 PM) kshepherd: tdonohue: ok cool, i'll let you know when i have a demo

(3:24:10 PM) ksclarke left the room (quit: Quit: Leaving.).

(3:34:49 PM) scottatm left the room (quit: *.net *.split).

(3:36:26 PM) jefftrimble left the room (quit: Quit: Leaving.).

(3:37:08 PM) PeterDietz: kshepherd, thanks for the pastebin. All, mea culpa for not being around. 1.7.RC2 will happen on Friday (perhaps into weekend if work is too busy)

(3:37:31 PM) scottatm [~scottatm@adhcp136.evans.tamu.edu] entered the room.

(3:38:56 PM) PeterDietz: Would have been nice to have had a git discussion, but it might be good for people to learn git first, or we could handle a discussion over current situation, future plans, ...