

CephFS - Cleanup #814

hadoop: refactor hadoop shim in terms of java libceph bindings

02/17/2011 07:14 AM - Sage Weil

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Normal		
<b>Assignee:</b>			
<b>Category:</b>			
<b>Target version:</b>	v0.54b		
<b>Tags:</b>		<b>Component(FS):</b>	Hadoop/Java
<b>Backport:</b>		<b>Labels (FS):</b>	Java/Hadoop
<b>Reviewed:</b>		<b>Pull request ID:</b>	
<b>Affected Versions:</b>			
<b>Description</b>			
Refactor the hadoop code in terms of generic Java bindings for libceph, instead of mixing the SWIG crap in with the actual functionality.			

History

#1 - 02/17/2011 09:00 AM - Greg Farnum

I'm not sure what you're after here -- you mean you want Java bindings for librados, and then the Hadoop patches should be updated to use those?

It's been a while, but from what I recall I was deliberately pretty careful to place Ceph functionality in the C code as much as possible so that we could update the interface and stuff without updating the Hadoop code (under the hope that it would eventually get merged into their codebase where we'd have more trouble getting at it). I suspect that maintaining the CephFS' ambivalence about how anything in libceph is actually used is something we want to preserve.

#2 - 02/17/2011 11:10 AM - Sage Weil

- I mean that, at the end of the day, we should probably have:
- a java Ceph interface binding that is identical to the interface exposed by libceph (something that looks more or less like posix, with whatever changes/additions make sense from a generic utility point of view). This should be useful to any Java thing that wants to interact with a ceph file system (not librados).
  - a hadoop class that uses the generic java Ceph bindings, and any hadoop-specific weirdness.

#3 - 02/18/2011 03:40 PM - Noah Watkins

Sage,

I see the general benefits of a Java-based libceph interface, but what are your long-term plans for Hadoop over Ceph? This sounds like there is room for a version that is not built on top of the kernel client, but the kernel client version is quite stable, and the tiny shim for locality information is fully integrated with the existing Hadoop native-code library.

Thanks,  
Noah

**#4 - 03/07/2011 10:37 AM - Sage Weil**

- translation missing: *en.field\_position* set to 531

**#5 - 03/07/2011 10:42 AM - Sage Weil**

- Target version changed from v0.26 to 12

- translation missing: *en.field\_position* deleted (531)

- translation missing: *en.field\_position* set to 10

**#6 - 04/29/2011 09:43 AM - Sage Weil**

- translation missing: *en.field\_position* deleted (44)

- translation missing: *en.field\_position* set to 1

- translation missing: *en.field\_position* changed from 1 to 644

**#7 - 06/13/2011 02:38 PM - Sage Weil**

- Target version changed from 12 to 19

**#8 - 08/31/2011 06:39 PM - Sage Weil**

- Target version changed from 19 to v0.37

**#9 - 08/31/2011 06:39 PM - Sage Weil**

- Target version changed from v0.37 to v0.36

**#10 - 08/31/2011 10:19 PM - Sage Weil**

- Target version changed from v0.36 to v0.37

- translation missing: *en.field\_position* deleted (762)

- translation missing: *en.field\_position* set to 1

**#11 - 08/31/2011 10:26 PM - Sage Weil**

- Target version deleted (v0.37)

- translation missing: *en.field\_position* deleted (4)

- translation missing: *en.field\_position* set to 2

**#12 - 09/02/2011 03:43 PM - Sage Weil**

- translation missing: *en.field\_position* deleted (12)

- translation missing: *en.field\_position* set to 34

**#13 - 09/02/2011 03:43 PM - Sage Weil**

- translation missing: *en.field\_position* deleted (34)

- translation missing: *en.field\_position* set to 14

**#14 - 11/29/2011 05:07 PM - Sage Weil**

<http://www.debian.org/doc/packaging-manuals/java-policy/x105.html>

**#15 - 05/30/2012 10:17 PM - Sage Weil**

- Status changed from New to In Progress

- translation missing: *en.field\_position* deleted (336)

- translation missing: *en.field\_position* set to 335

**#16 - 10/08/2012 05:33 PM - Sage Weil**

- *Project changed from Ceph to CephFS*
- *Category deleted (20)*

**#17 - 10/08/2012 05:34 PM - Sage Weil**

- *translation missing: en.field\_position deleted (463)*
- *translation missing: en.field\_position set to 1*

**#18 - 10/08/2012 05:35 PM - Sage Weil**

- *Category set to 48*

**#19 - 10/09/2012 01:32 PM - Sage Weil**

- *Target version set to v0.54b*
- *translation missing: en.field\_position deleted (1)*
- *translation missing: en.field\_position set to 3*

**#20 - 12/24/2012 02:58 PM - Sage Weil**

- *Status changed from In Progress to Resolved*
- *translation missing: en.field\_position deleted (12)*
- *translation missing: en.field\_position set to 12*

**#21 - 07/13/2016 12:31 AM - Greg Farnum**

- *Component(FS) Hadoop/Java added*

**#22 - 03/09/2019 12:29 AM - Patrick Donnelly**

- *Category deleted (48)*
- *Labels (FS) Java/Hadoop added*