

Linux kernel client - Bug #2298

rbd: broken encode_op for big-endian hosts?

04/13/2012 10:40 PM - Sage Weil

Status:	Resolved	Start date:	04/13/2012
Priority:	High	Due date:	
Assignee:	Alex Elder	% Done:	0%
Category:	rbd	Estimated time:	0.00 hour
Target version:	v3.4	Spent time:	0.00 hour
Source:	Community (dev)	Severity:	3 - minor
Tags:		Reviewed:	
Backport:		Affected Versions:	
Regression:	No	ceph-qa-suite:	

Description

Date: Sat, 14 Apr 2012 03:34:20 +0100
From: Al Viro <viro@ZenIV.linux.org.uk>
To: Sage Weil <sage@newdream.net>
Cc: linux-kernel@vger.kernel.org, linux-fsdevel@vger.kernel.org
Subject: osd_req_encode_op() breakage?

```
static void osd_req_encode_op(struct ceph_osd_request *req,  
                             struct ceph_osd_op *dst,  
                             struct ceph_osd_req_op *src)
```

```
{  
    dst->op = cpu_to_le16(src->op);
```

```
    switch (dst->op) {  
    case CEPH_OSD_OP_READ:  
    case CEPH_OSD_OP_WRITE:
```

is an interesting thing to say, seeing that CEPH_OSD_OP_READ et.al. are all host-endian... Should that be "switch (src->op)" instead? AFAICS, that sucker had appeared in that form back in commit 68b4476b0bc13fef18266b4140309a30e86739d2

Author: Yehuda Sadeh <yehuda@hq.newdream.net>
Date: Tue Apr 6 15:01:27 2010 -0700

ceph: messenger and osdc changes for rbd

and it seems to be broken on big-endian hosts. Doesn't look like a misspelled le16_to_cpu() either, since dst->op ends up going on the wire...

I'm really mystified by that - it looks like it must've shown up immediately on big-endian hosts; it's not like it was an obscure codepath, after all...

Comments?

--

To unsubscribe from this list: send the line "unsubscribe linux-fsdevel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

History

#1 - 04/16/2012 01:46 PM - Alex Elder

I haven't looked at this in any detail but I presume Al is correct. We don't have any big endian hardware anywhere, do we? Might be worth finding an old PowerPC or MIPS machine for occasional tests to verify these sorts of things.

#2 - 04/16/2012 02:15 PM - Sage Weil

there are some old g5's in the closet here at aon that we can use.

in the past we've found/fixed these issues with sparse.. yehuda runs it periodically and it picks up the __le32 vs u32 eerrors.

#3 - 04/18/2012 08:22 AM - Alex Elder

- *Status changed from New to In Progress*

- *Assignee set to Alex Elder*

I sent a note to the various lists Al Viro posted to, to confirm the bug (wasn't sure whether Sage had or not).

I also implemented his suggested fix and posted that to the lists and will be running it through tests as soon as it's built.

We'll want this one to go to Linus for the current release cycle (3.4).

I also promised to ensure we were going to be doing endianness testing for our software on a regular basis.

#4 - 04/20/2012 07:12 PM - Alex Elder

- *Status changed from In Progress to Resolved*

This has been fixed. I have been testing it in a private branch and will shortly be updating the ceph-client testing to have it.

commit d8bc680e64a856ed2489f2df89e9561bc3aa70b2

Author: Alex Elder <elder@dreawmhost.com>

Date: Fri Apr 20 15:49:43 2012 -0500

```
ceph: osd_client: fix endianness bug in osd_req_encode_op()
```