# Kerberos and security with Lustre

**Sebastien Buisson**
Parallel File Systems
BDS R&D Data Operations
sebastien.buisson@atos.net

# Kerberos & security

- Kerberos related patches
  - patches role in GSS/Kerberos architecture
- Remaining work
  - enc_pools issue
  - documentation
  - cross-realm authentication

# Lustre/GSS Software stack

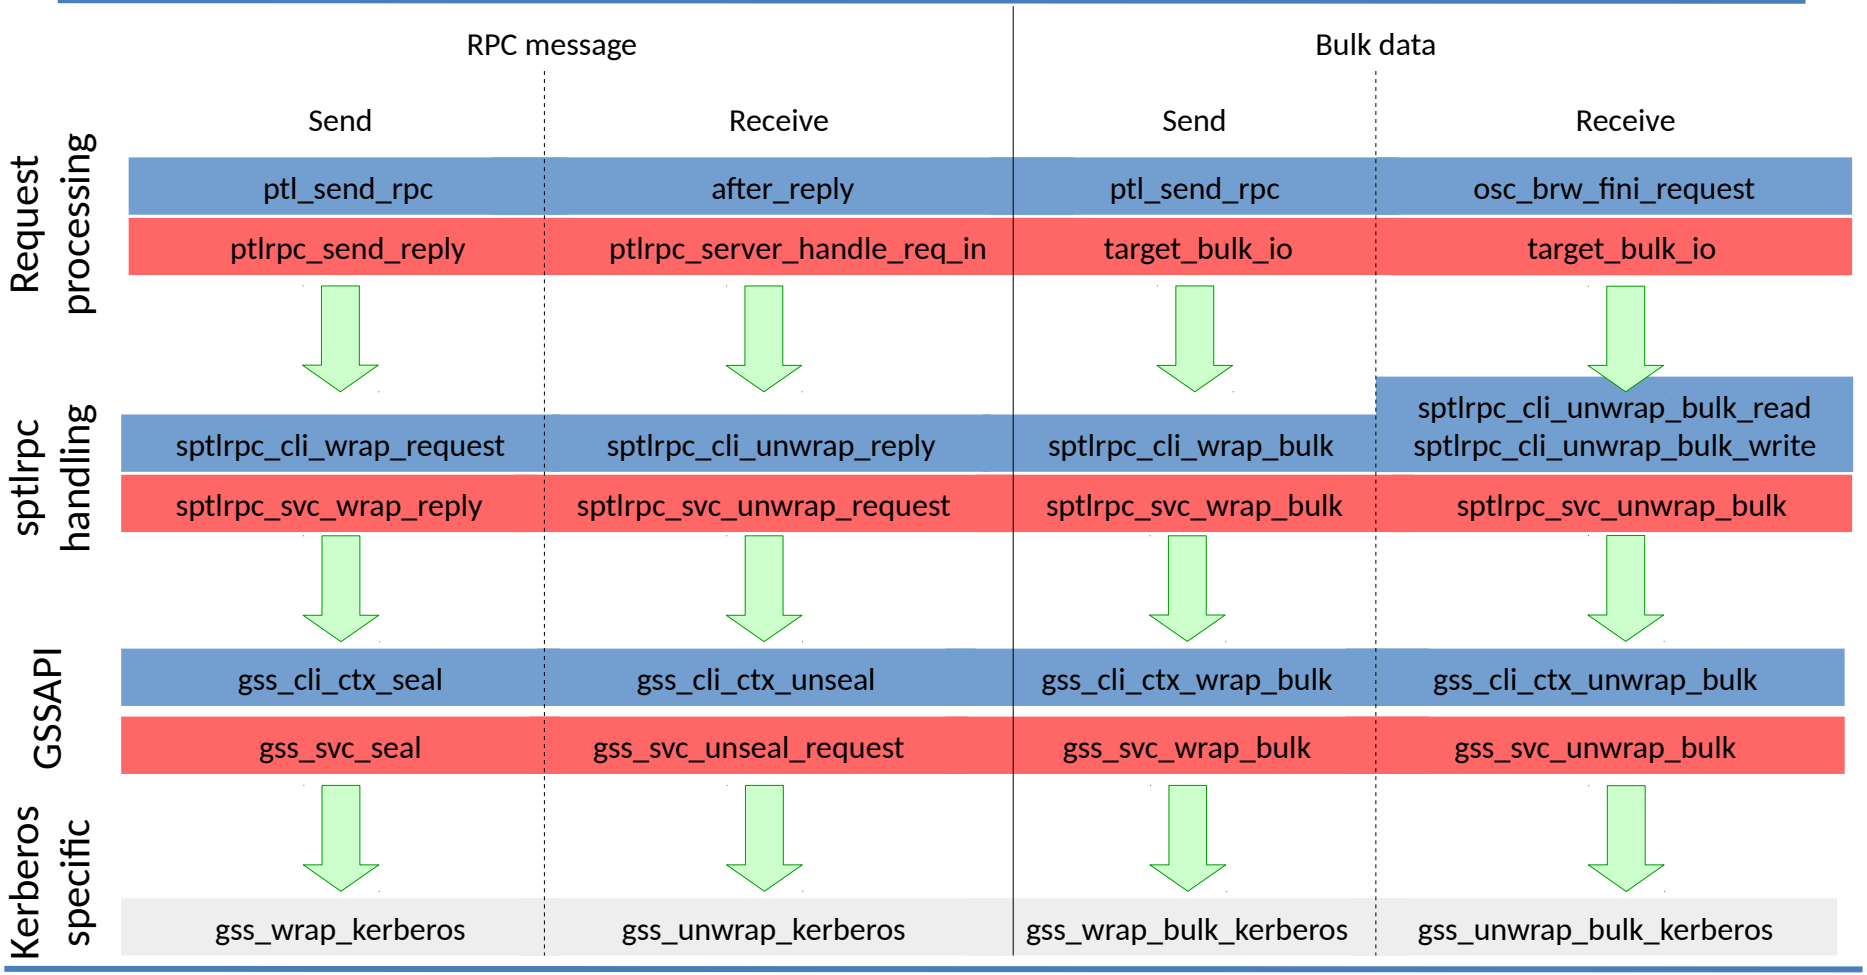

# Kerberos related patches

- LU-3778 (http://review.whamcloud.com/14040)
  - sptlrpc subsystem initialized for OSP and LWP connections
    => call to sptlrpc_lprocfs_cliobd_attach()
  - prevent LBUG in GSS related functions when dealing with OSP and LWP OBDs

- Profitable to sptlrpc layer

# Kerberos related patches

- LU-6356 (http://review.whamcloud.com/14041)
- LU-6356 (http://review.whamcloud.com/14043)
  - handle sec context requests properly

- Profitable to security context processing

# Kerberos related patches

- LU-6356 (http://review.whamcloud.com/14349)
  - fix security flavor setting for connection to mgs

- Profitable for every secured connection to MGS

# Kerberos related patches

- LU-6356 (http://review.whamcloud.com/14042)
  - call out info must include 'self nid'

- Profitable to Kerberos principal checking, but not limited to this usage

# Kerberos related patches

- LU-6020 (http://review.whamcloud.com/14018)
  - fix bulk nob on bulk writes
- LU-6020 (http://review.whamcloud.com/14020)
  - wrap readdir bulk replies

- Profitable to bulk wrap part of sptlrpc layer

# Kerberos related patches

- LU-6020 (http://review.whamcloud.com/13076)
  - fix misuse of krb5int_derive_key()
- LU-6020 (http://review.whamcloud.com/13631)
  - proper sg list initialization in krb5 digest/encrypt/decrypt functions

- Kerberos specific

# Kerberos related patches

- Patches not landed yet:
  - LU-3778 (http://review.whamcloud.com/14040)
  - LU-6356 (http://review.whamcloud.com/14349)
  - LU-6356 (http://review.whamcloud.com/14041)
  - LU-6356 (http://review.whamcloud.com/14042)
  - LU-6356 (http://review.whamcloud.com/14404)

- Please review :)
  - Is May target realistic (Lustre 2.8)?

# Kerberos & security

- Kerberos related patches
  - patches role in GSS/Kerberos architecture
- Remaining work
  - enc_pools issue
  - documentation
  - cross-realm authentication

# enc_pool issue

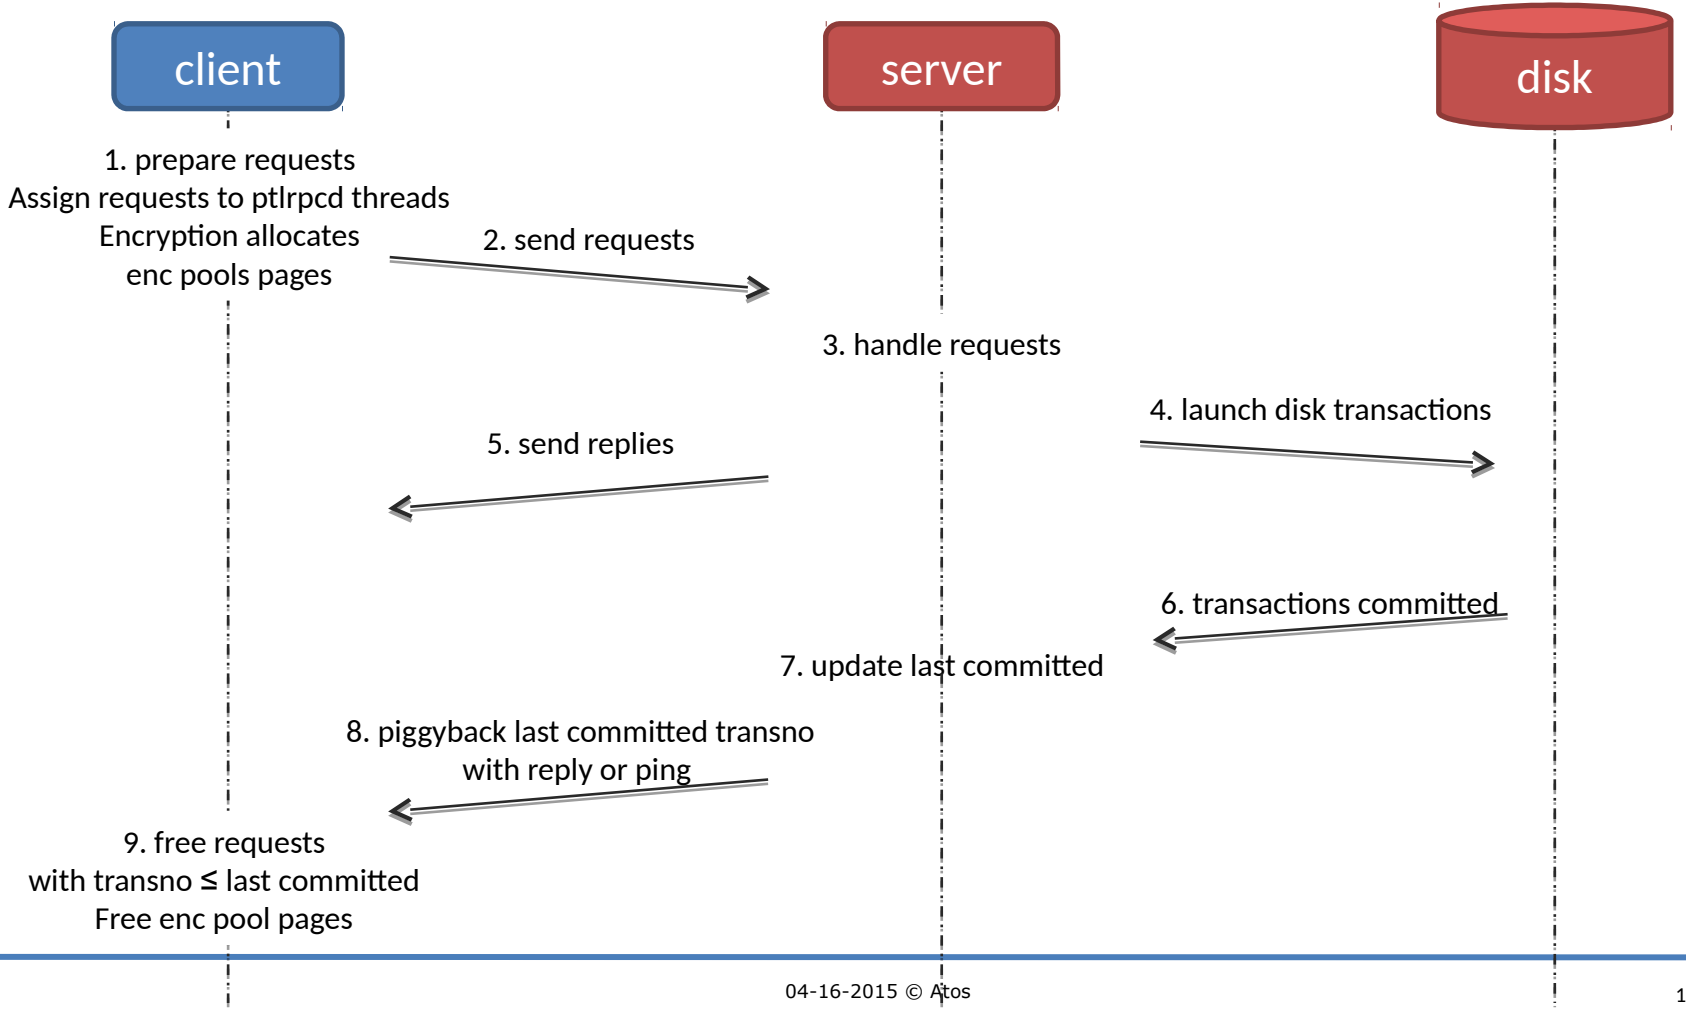

# enc_pool issue

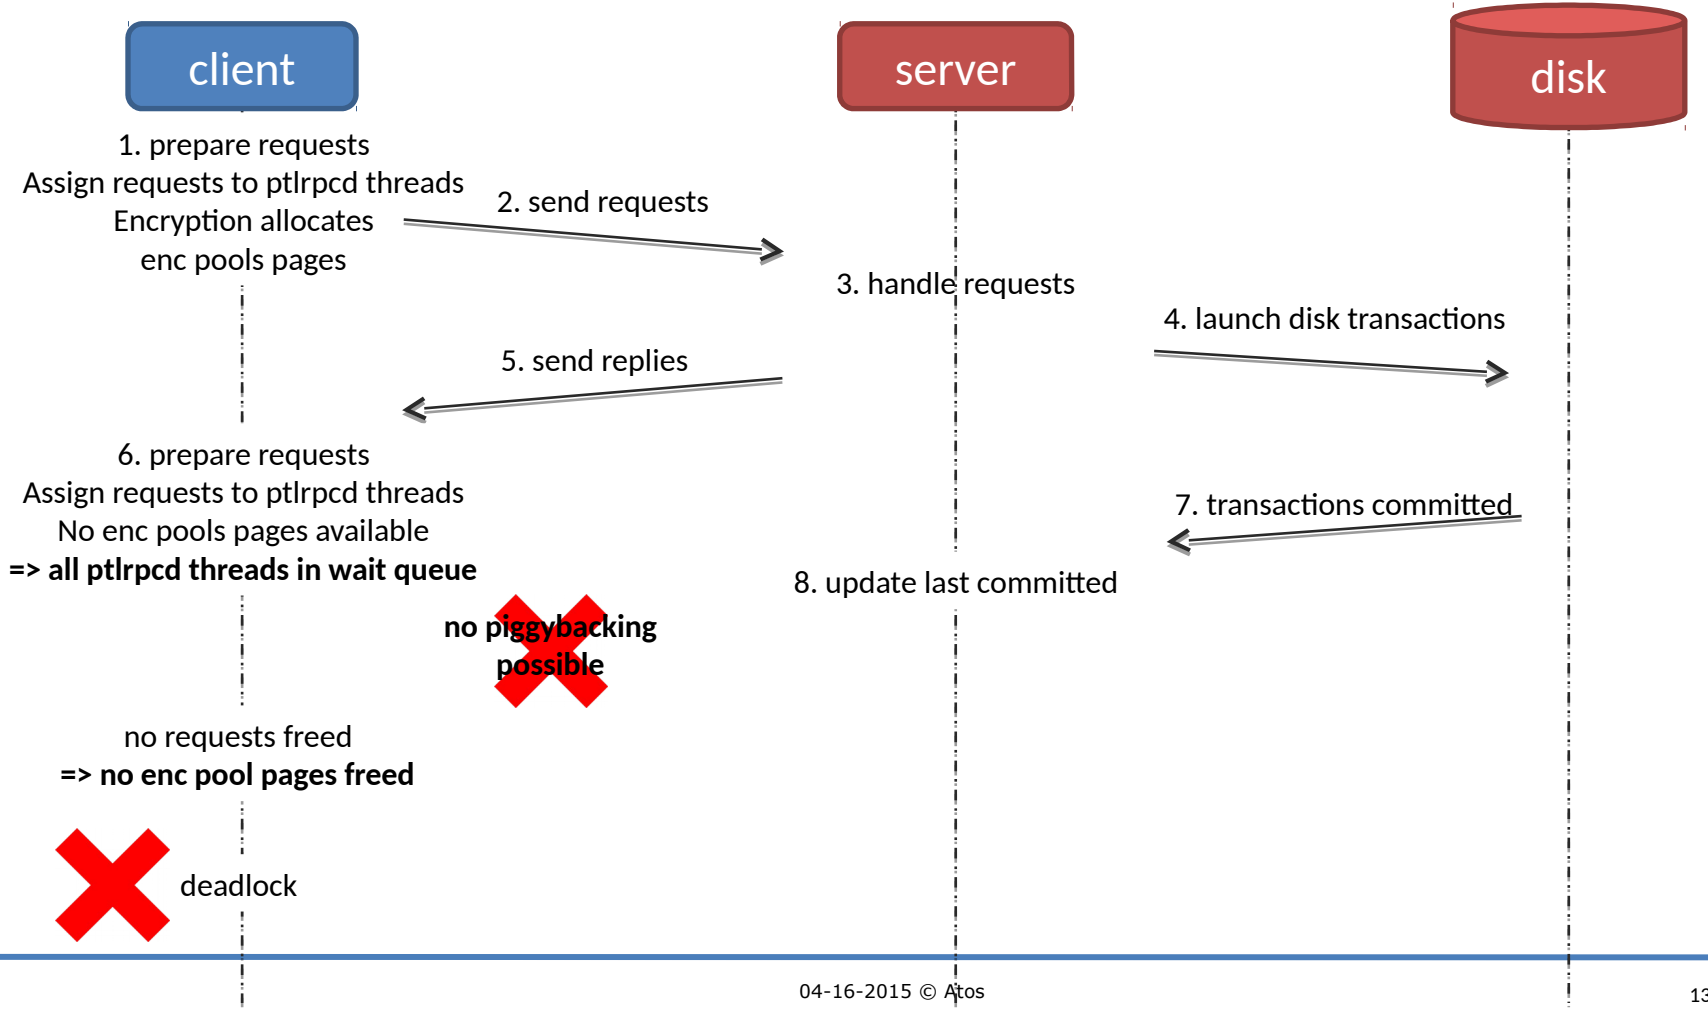

# enc_pool issue

- Possible solutions:

a) make epp_max_pages tunable via an sptlrpc kernel module parameter

b) free enc_pool pages as soon as request is sent

- adds latency to request processing
- ok with resend/replay?

c) prevent last thread to enter wait queue

=> instead, stop request processing and put back in request queue

d) try to allocate enc_pool pages before request is assigned to a ptlrpcd thread

=> if no availability, put back in request queue

# Documentation

- Lustre wiki
  - http://wiki.old.lustre.org/index.php/GSS_/_Kerberos
  - reintegrate in new wiki?
    - updates needed to reflect current code

# Cross-realm authentication

- Can idmap feature from IU permit Kerberos cross-realm authentication?

# Thanks

For more information please contact:
sebastien.buisson@atos.net



04-16-2015