

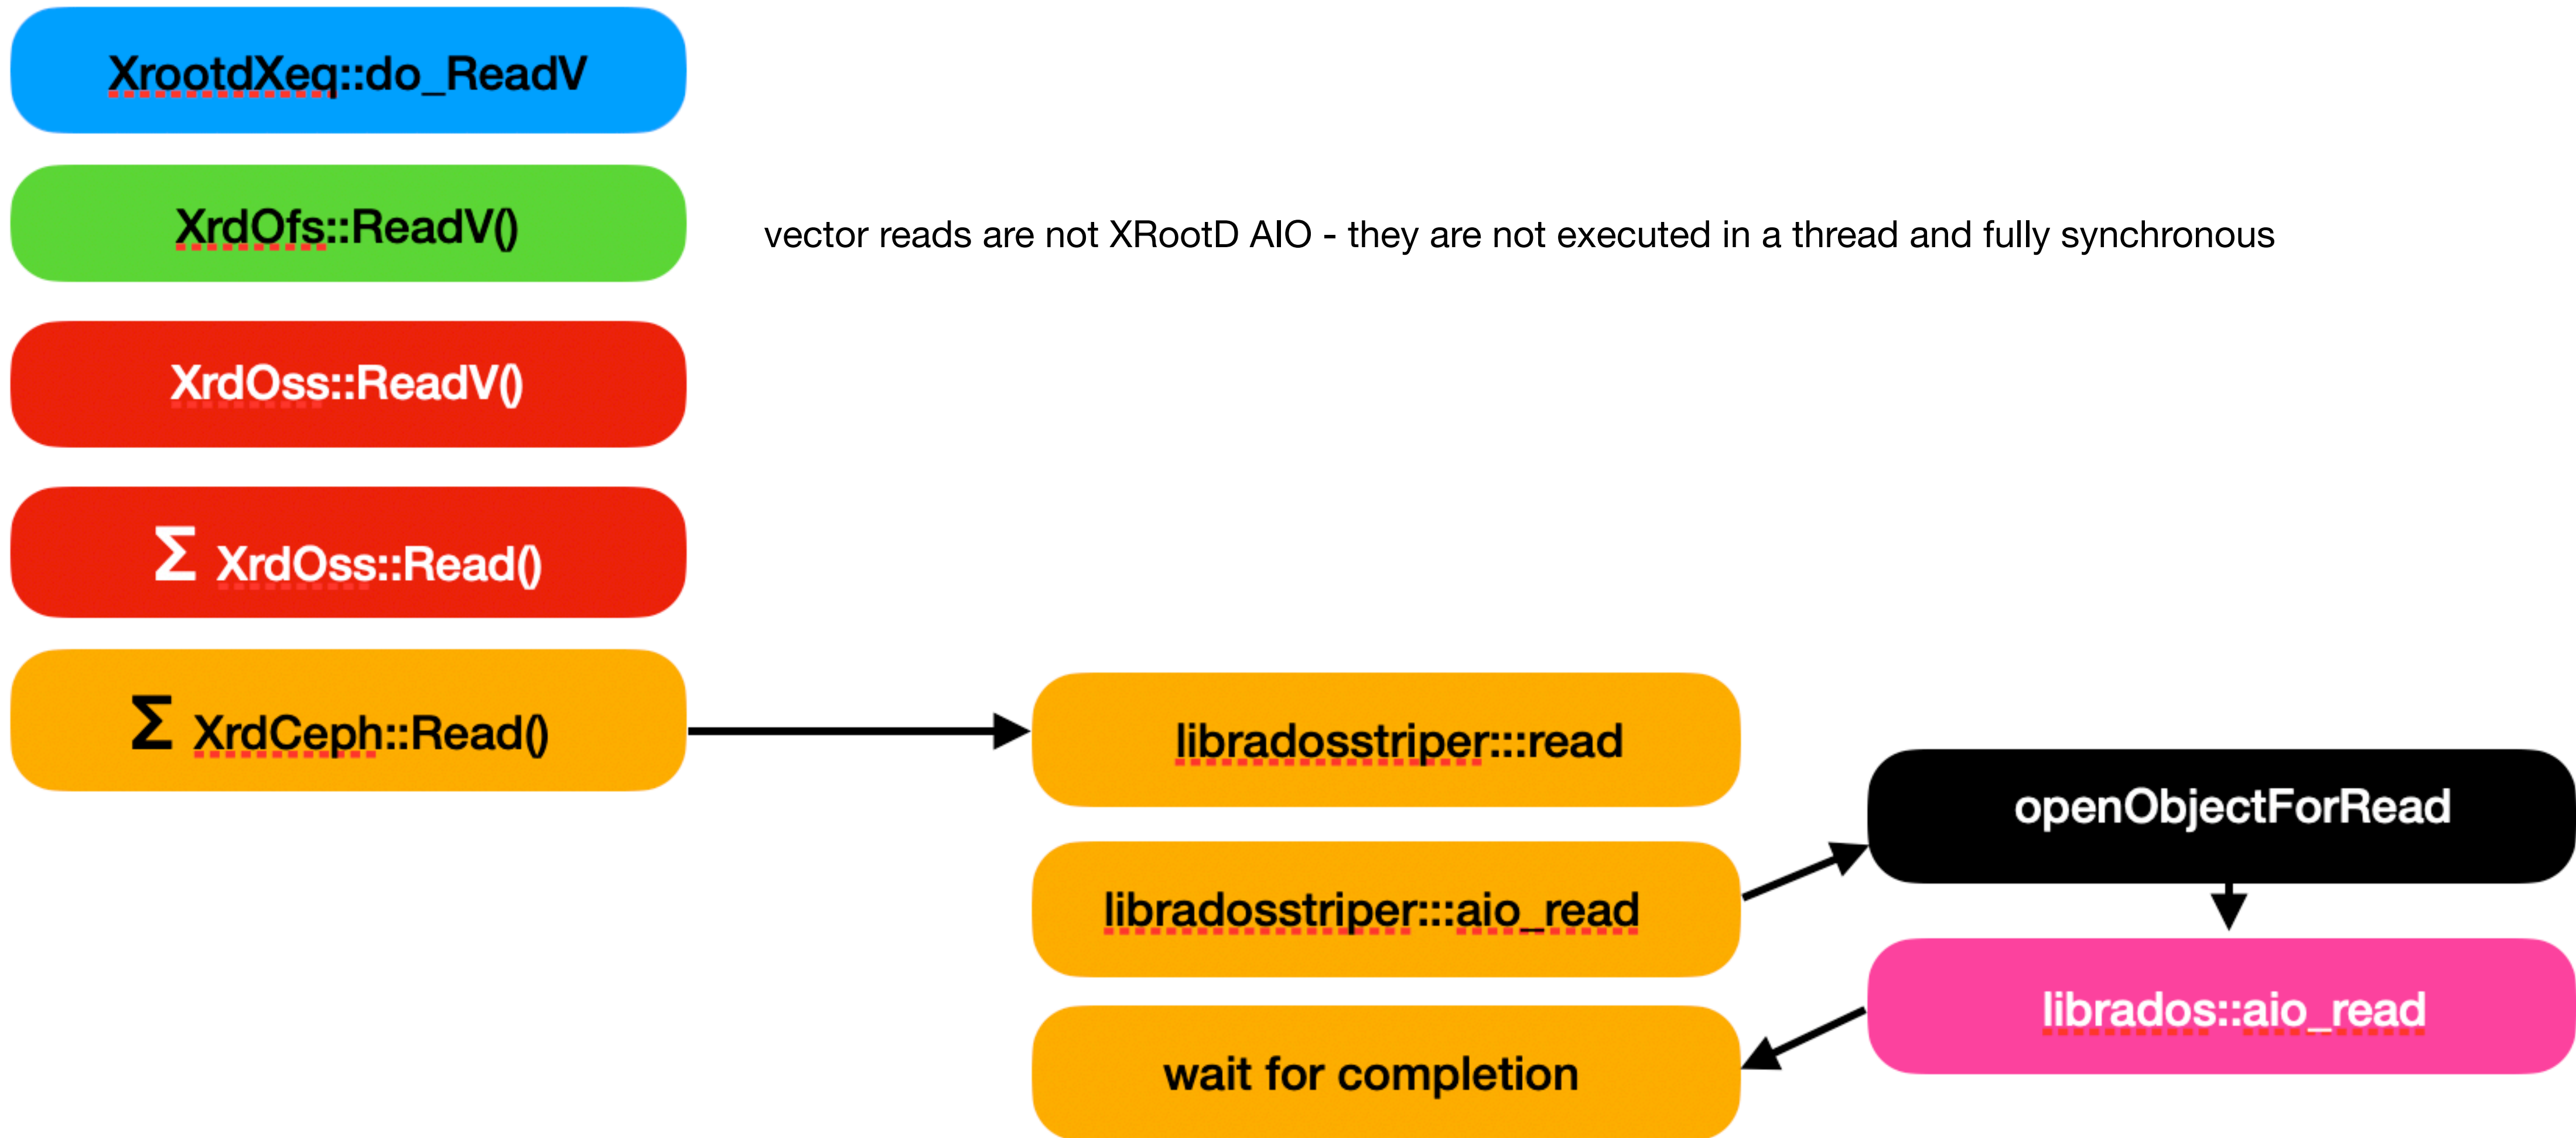
# **XrdCeph**

## **how-to get real async IO**

A.J.Peters

# readv call diagram

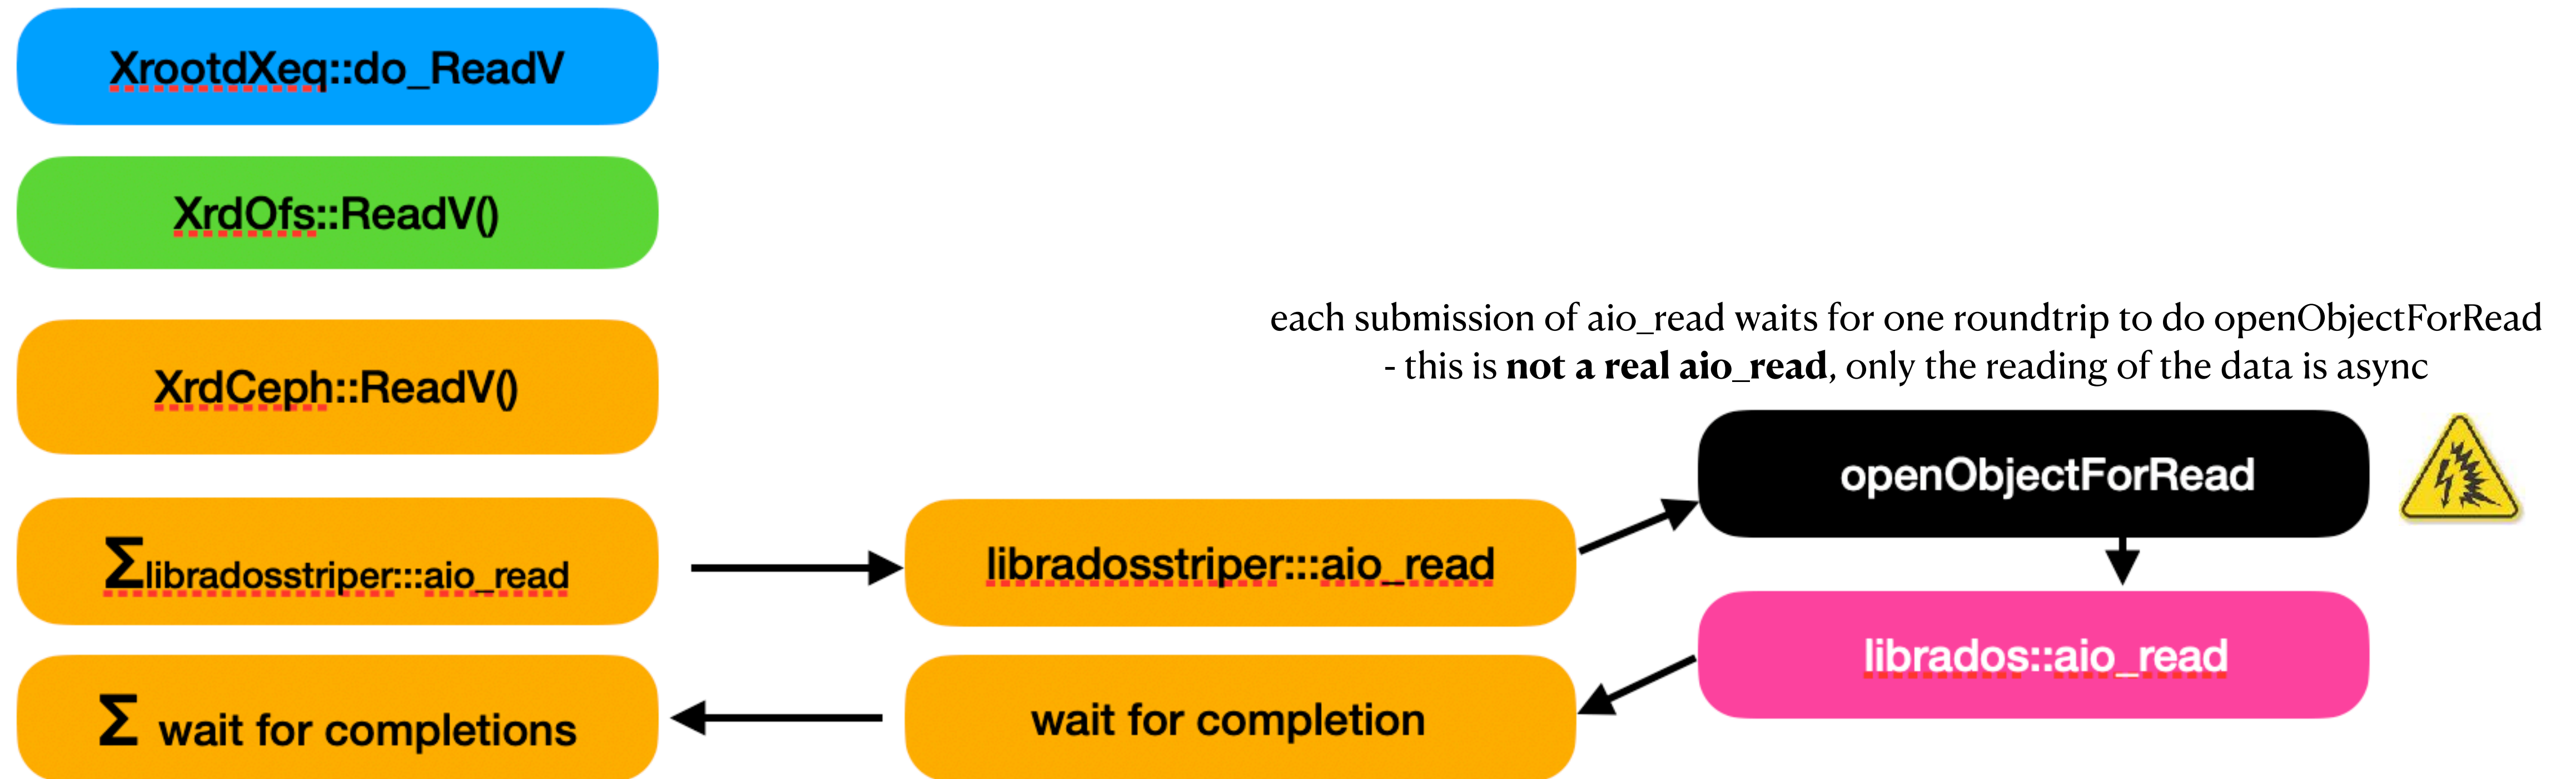
readv are not AIO - not executed in asynchronous threads



# readv call diagram

## first possible improvement

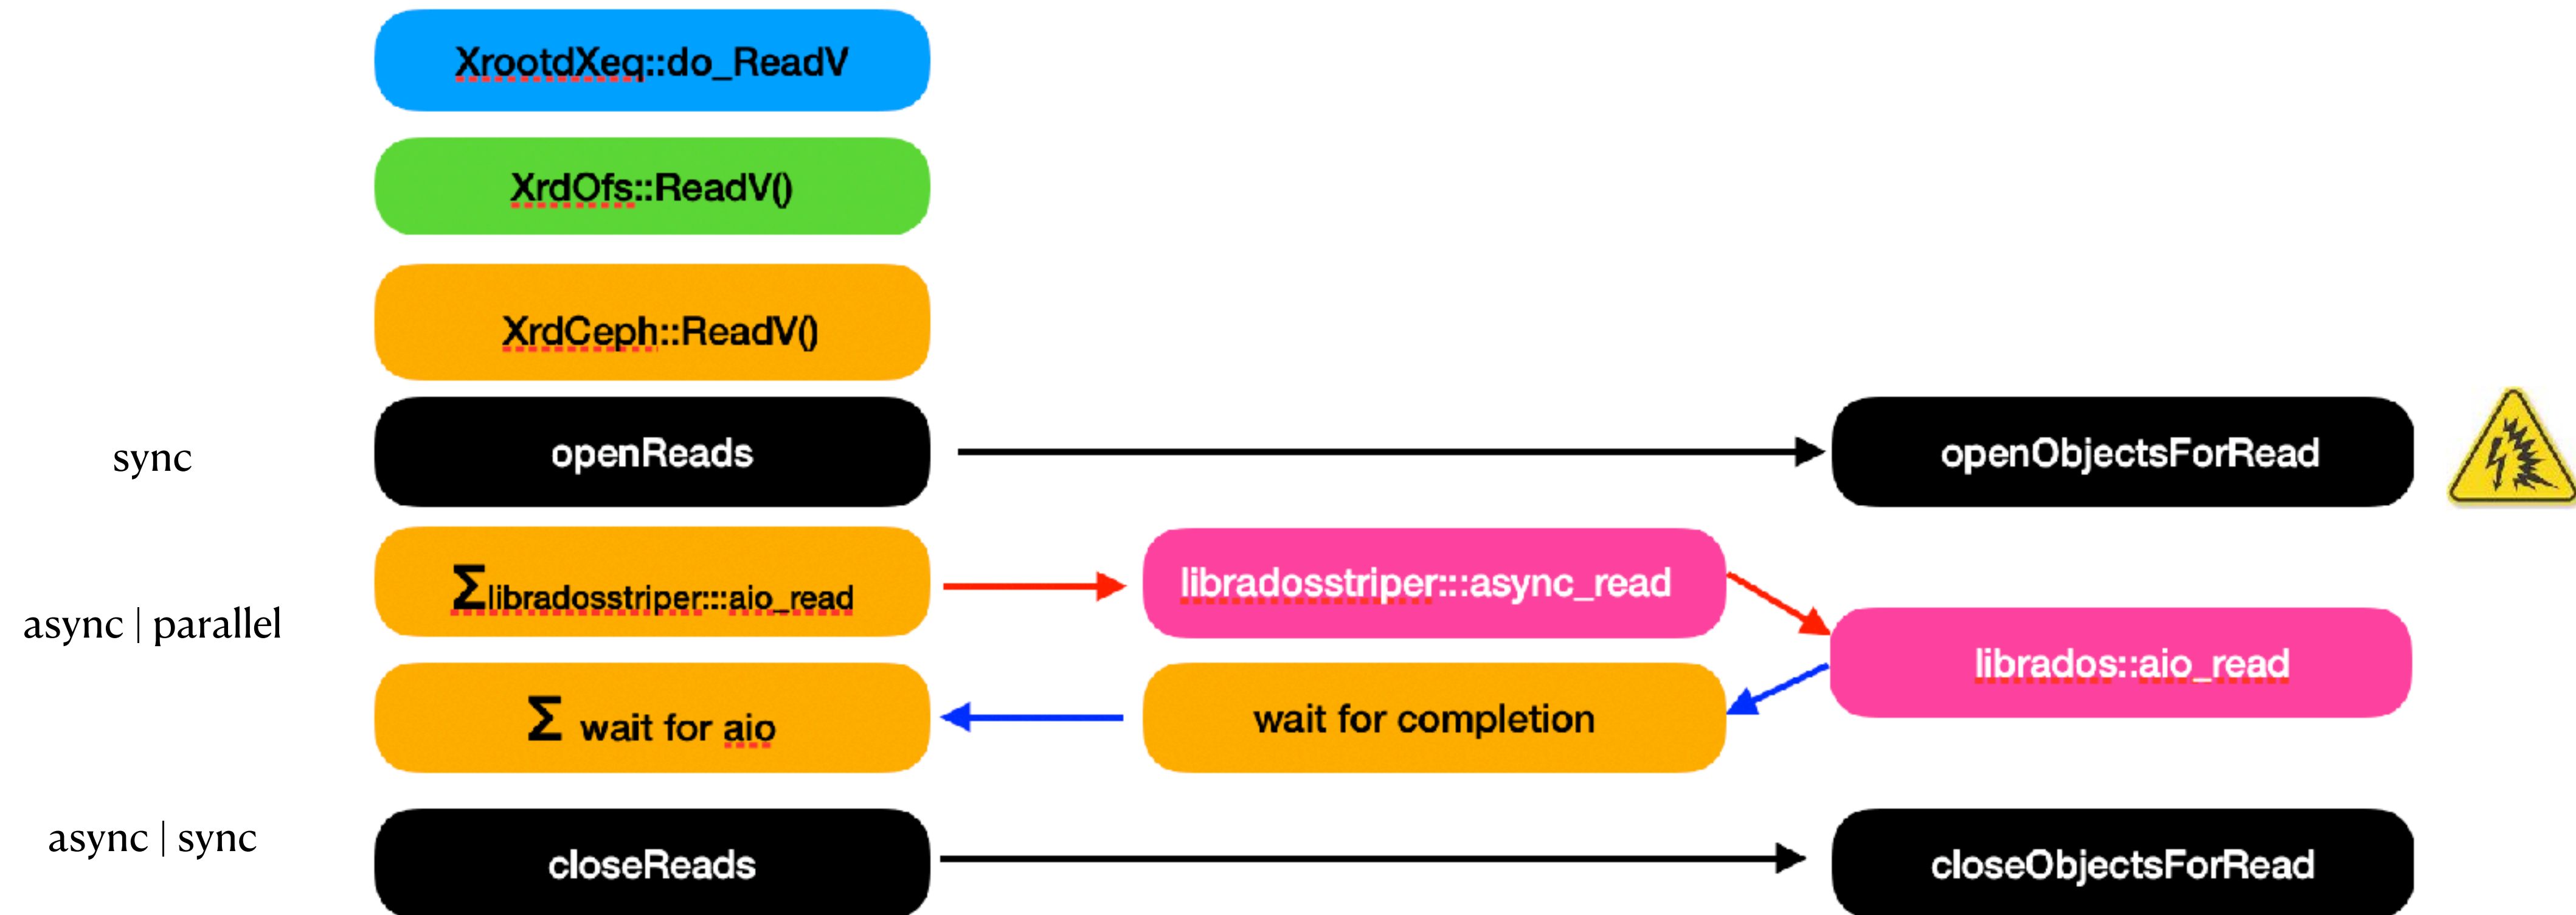
does not touch libradosstriper



# readv call diagram

best possible improvement

*does touch libradosstriper*



# general improvement

## bi-sect libradosstriper API

- to lock every read was intended to allow interleaved read-writes
  - de-facto completely useless functionality for **WORM**
    - this is an 'anti performance' implementation
  - much better split the api into:
    - **lock** for reading - tie to *do\_Open*
    - do many **aio\_reads** and **collect** them (*do\_Read do\_ReadV*)
    - **unlock** for reading - tie to *do\_Close*