

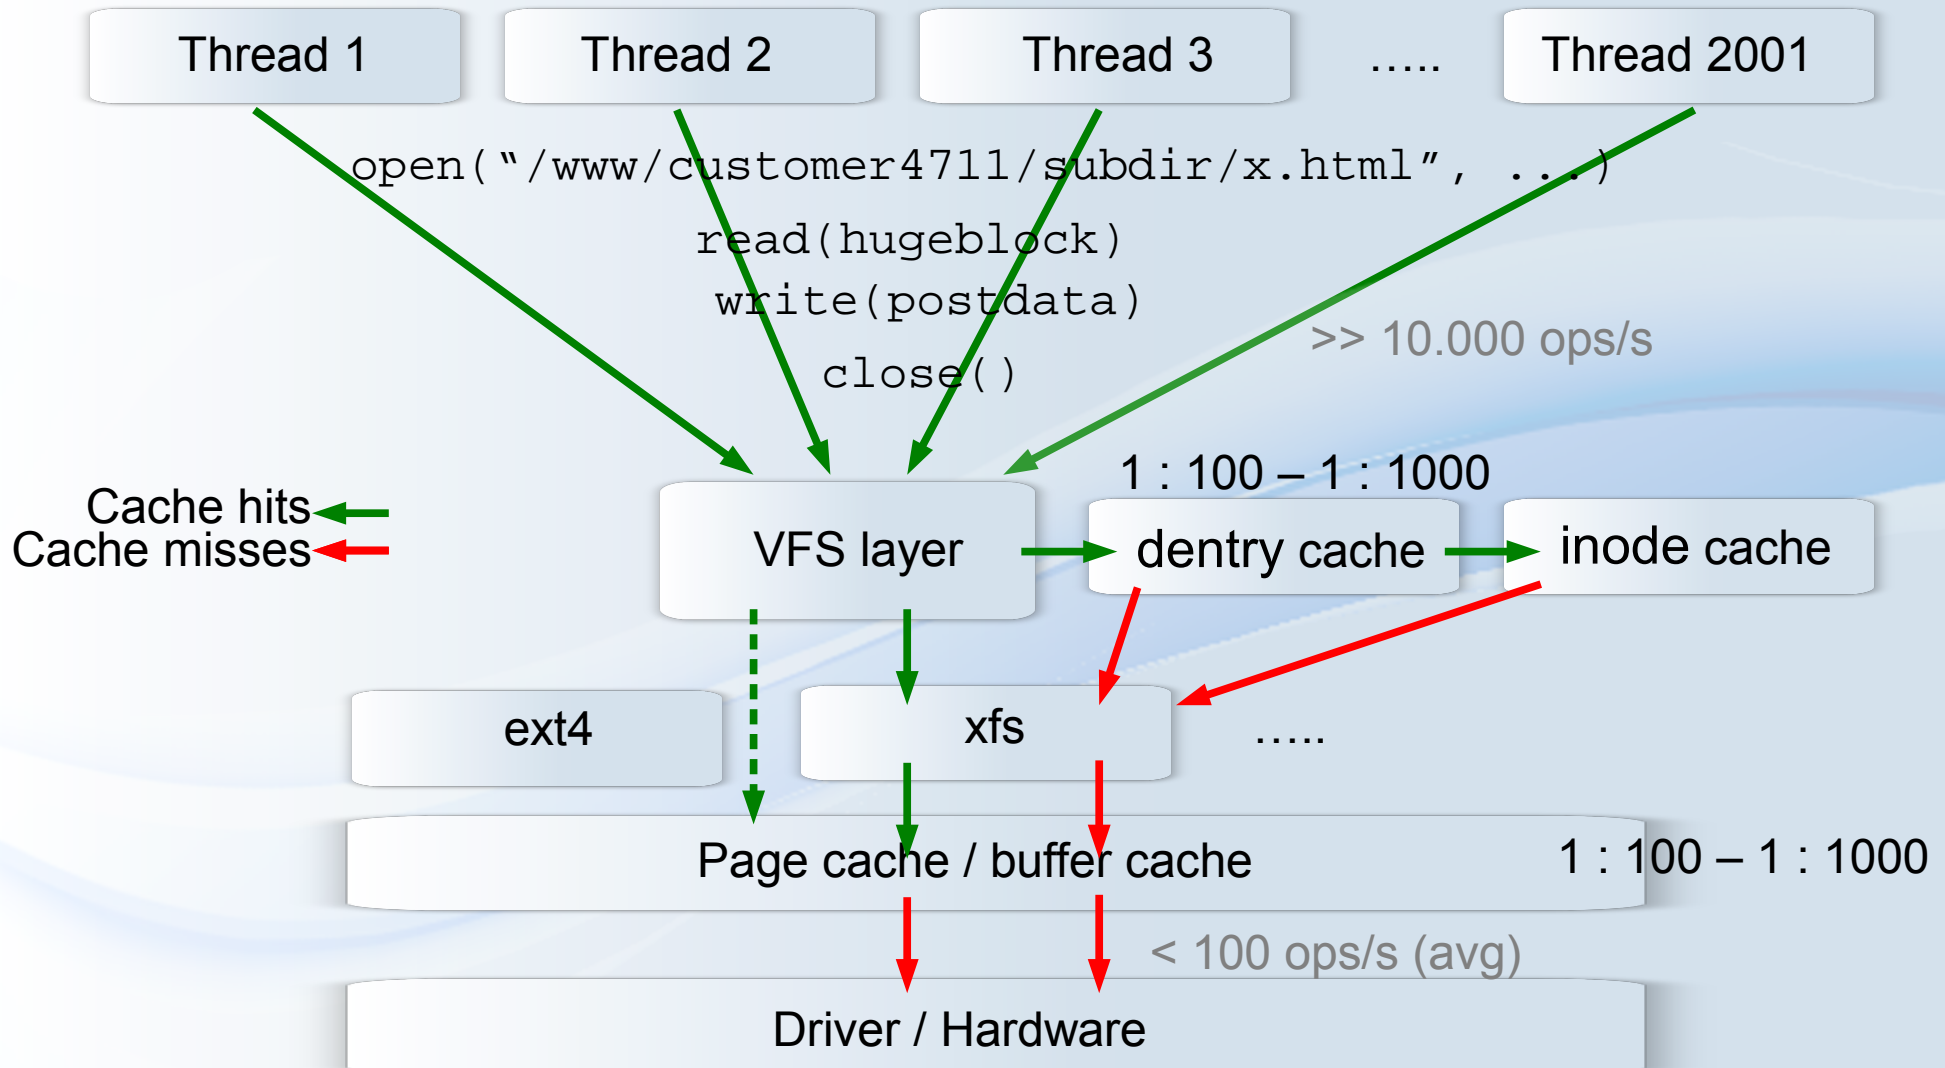
The Linux IO Stack unveiled



Short-Term Substitute Presentation at LinuxTAG 2013
by Thomas Schöbel-Theuer


- **Where does Linux performance come from?**
- **How does a cache work?**
- **Some tuning hints**
- **Details of IO stack by Werner Fischer & co**

Where does Linux performance come from?



How does a cache work?


Start:



After access to D:



Finally release of D:



aka **LRU** strategy: keep the HOTTEST items, remove the coldest first

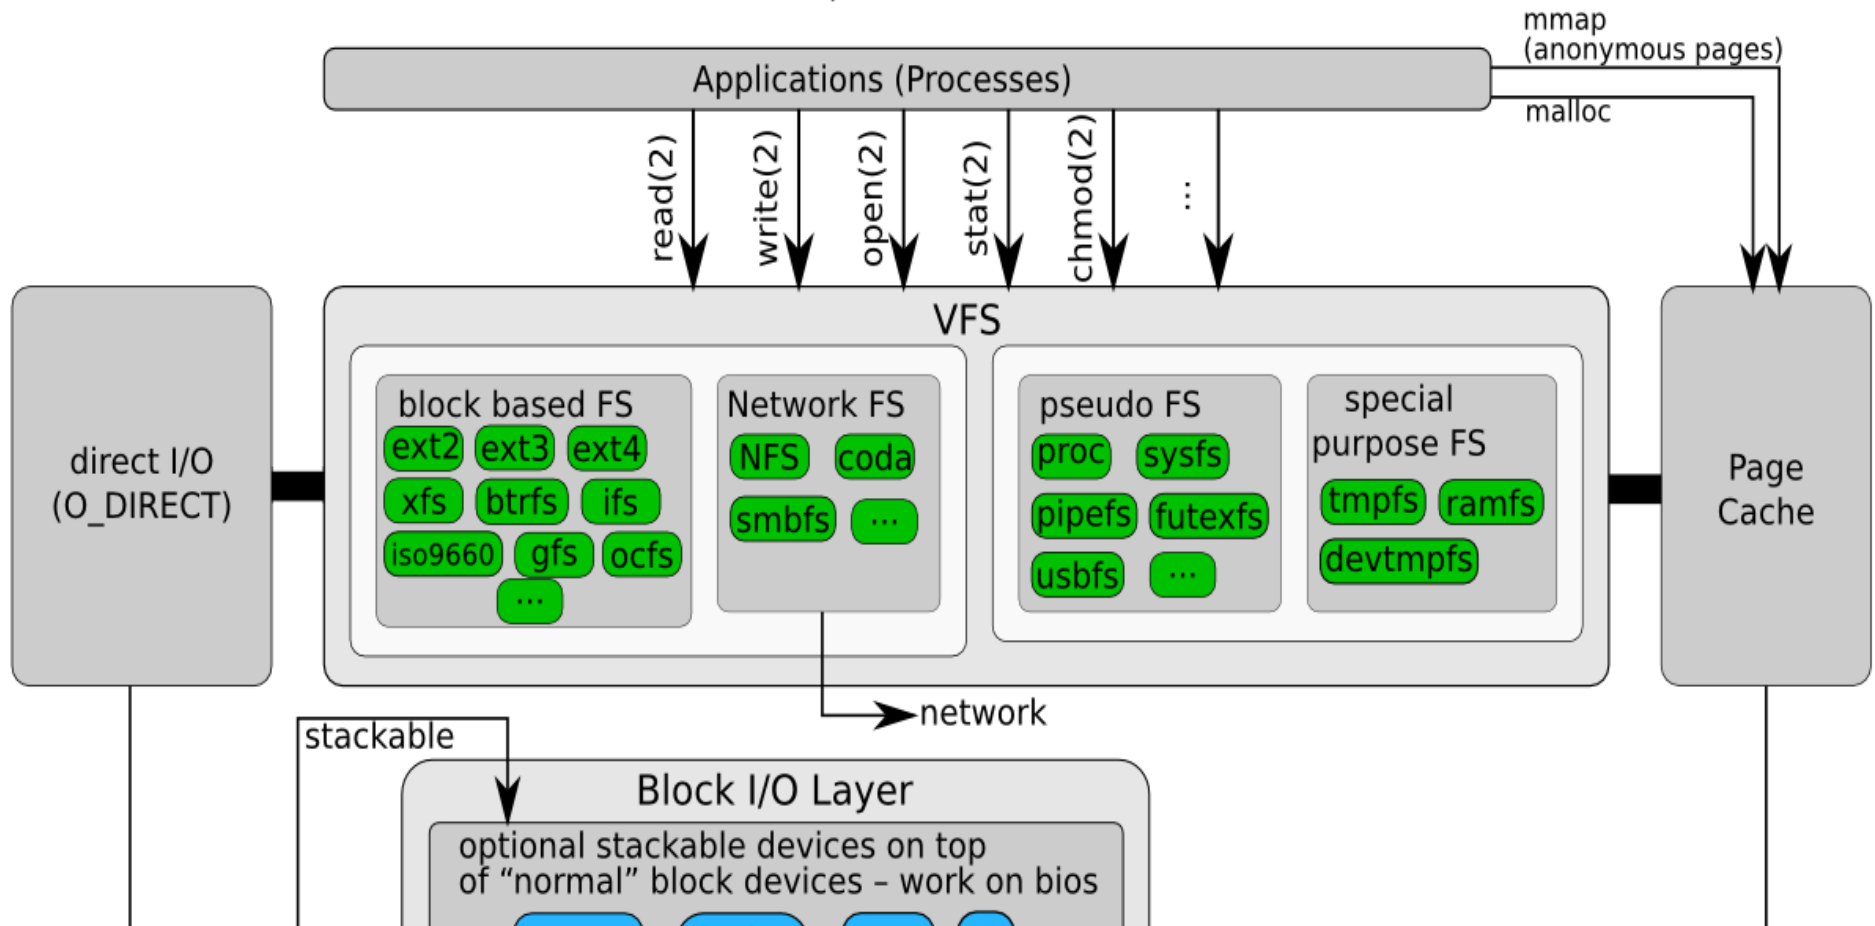
Some tuning hints

- **Before tuning anything else, tune your caches!**
 - look at `/proc/slabinfo`
 - don't buy more expensive hardware if you are not sure that you *really* need it
- **Theory of cache tuning: Denning ~1968**
 - aka Workingset Theory
 - description at blkreplay.org

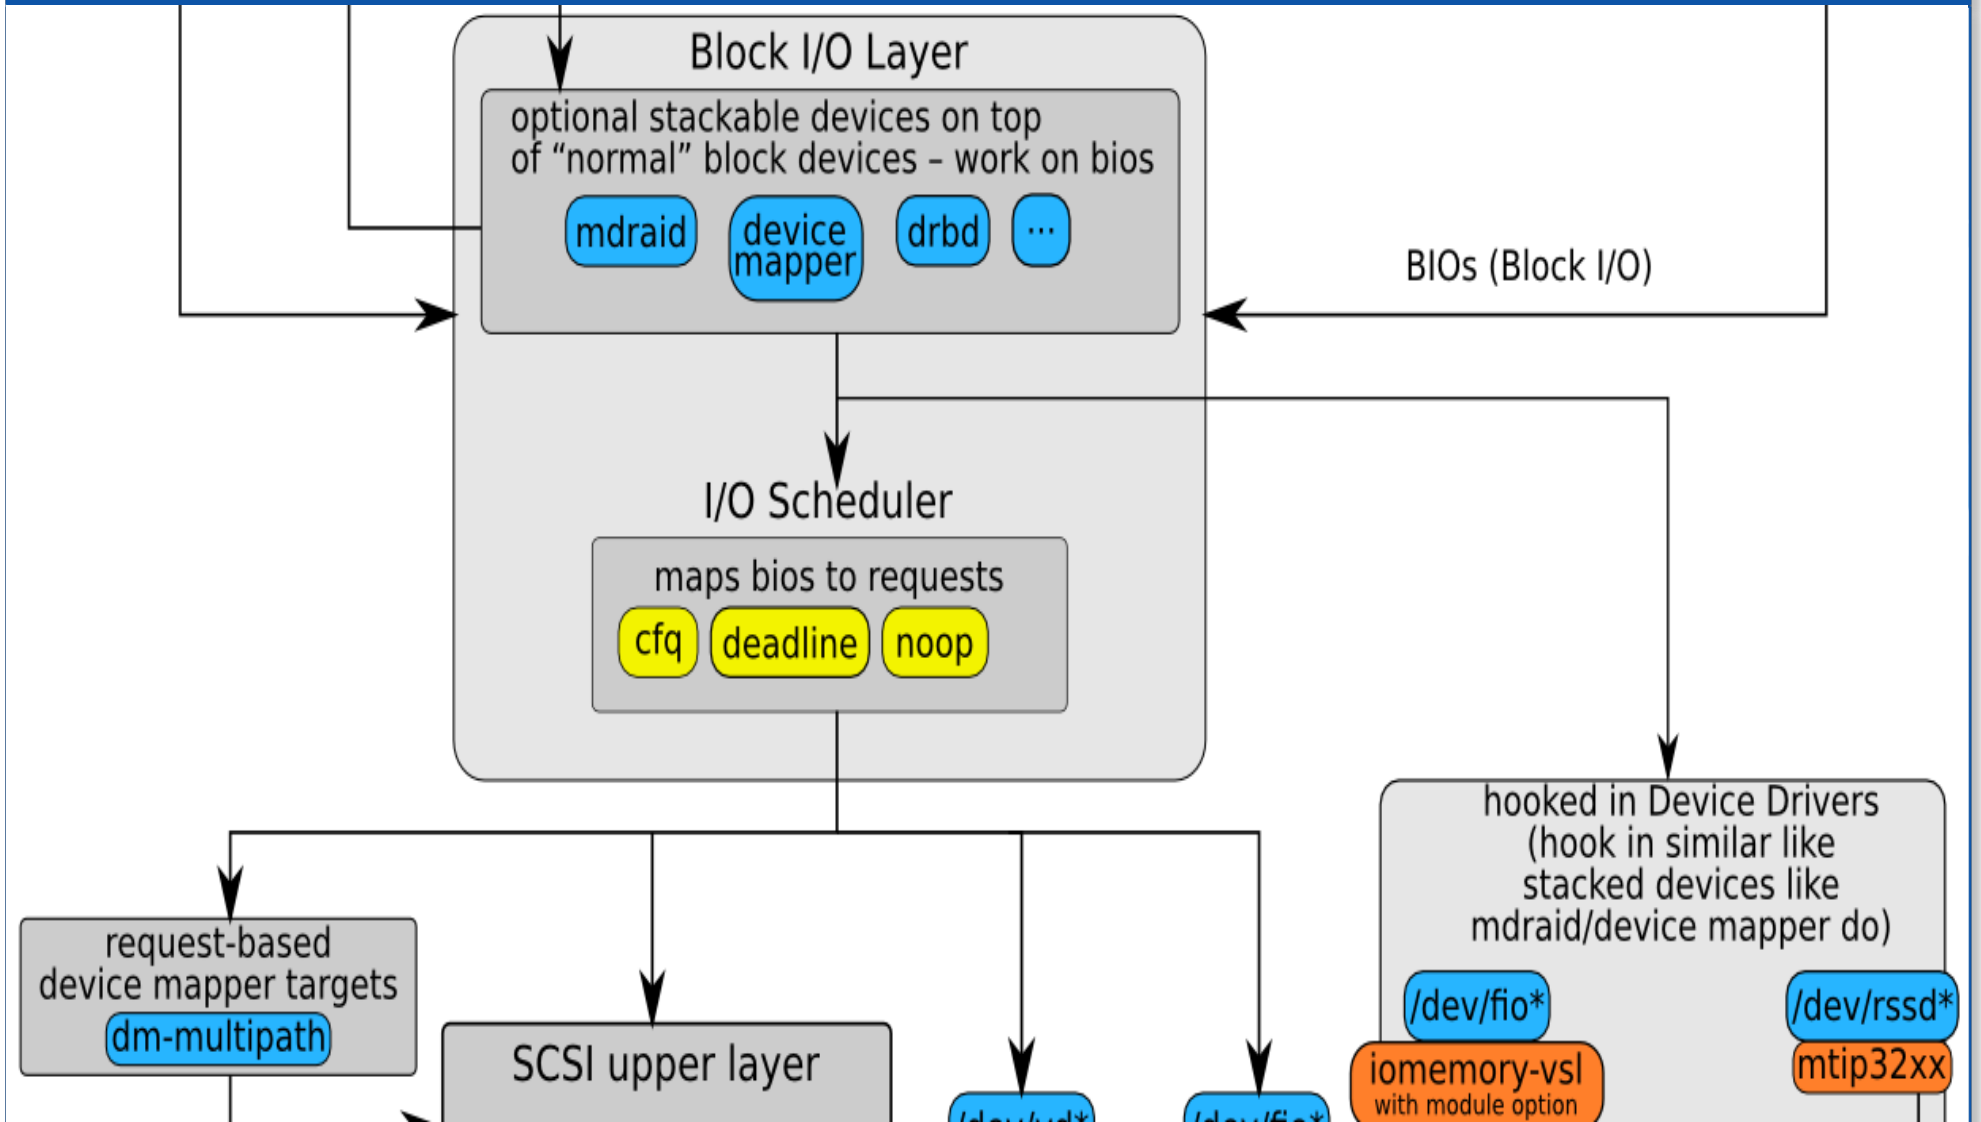


The Linux I/O Stack Diagram

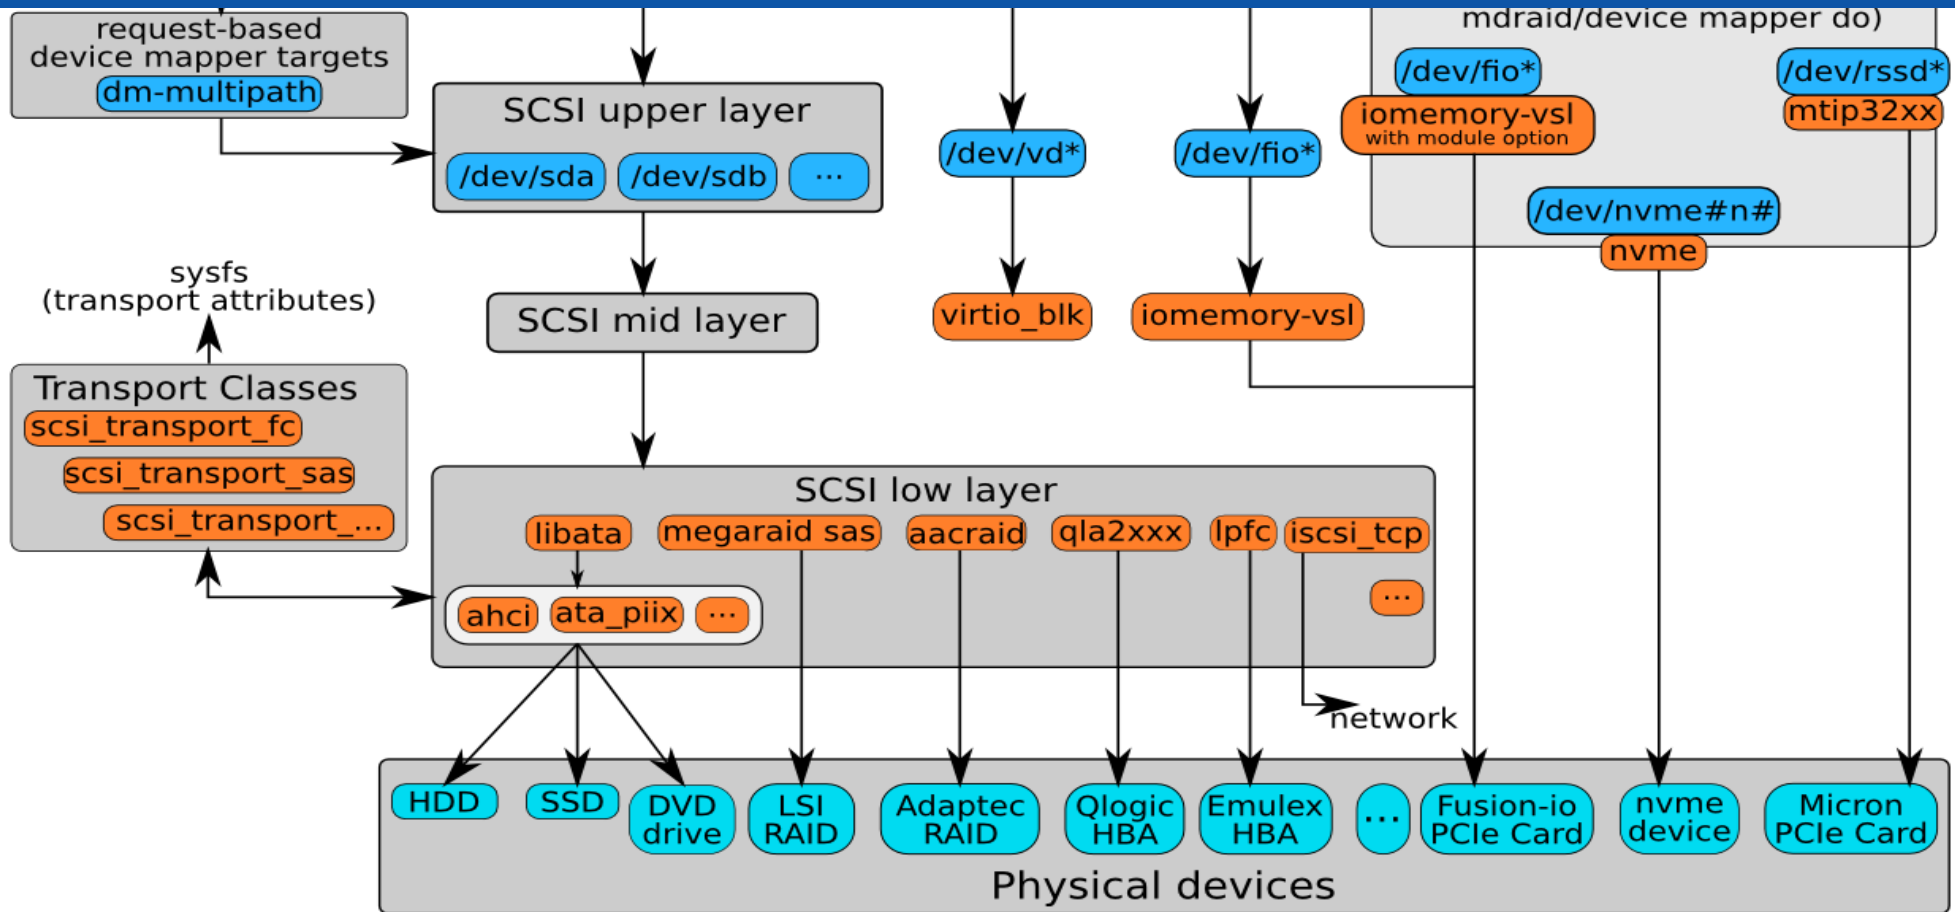
version 1.0, 2012-06-20
outlines the Linux I/O stack as of Kernel version 3.3



IO Stack Details (2) by Werner Fischer & co



IO Stack Details (3) by Werner Fischer & co



The Linux I/O Stack Diagram (version 1.0, 2012-06-20)
<http://www.thomas-krenn.com/en/oss/linux-io-stack-diagram.html>
Created by Werner Fischer and Georg Schönberger
License: CC-BY-SA 3.0, see <http://creativecommons.org/licenses/by-sa/3.0/>