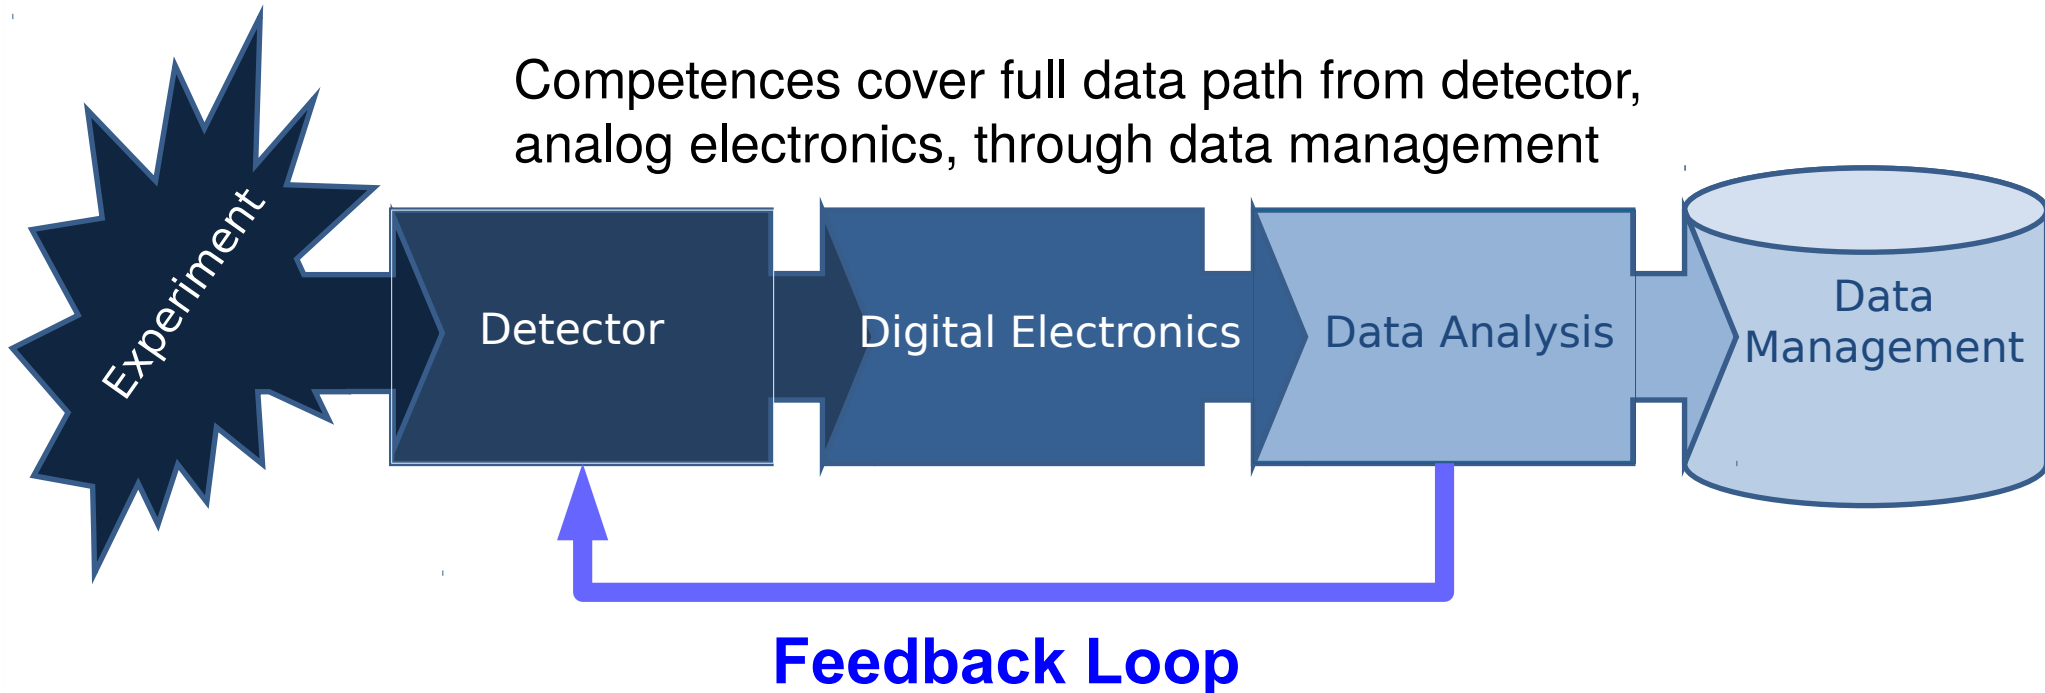



UFO - A platform for high data rate instrumentation with GPUs



- ▶ Programmable streaming DAQ platform that can be adopted easily to image sensors and fast sampling DACs
- ▶ Fully pipelined image-processing framework supporting diversity of the hardware platforms and based on open standards
- ▶ Integration in control systems, e.g. for image-based feedback

High Speed DAQ and Computing at IPE


Fast tomography with image-based feedback




X-ray cine-tomography




uPIV for thin films




Electron beam monitoring




Diffraction laminography for crystallography




Virtual lab for research of arthropod structures



ROFEX for fluid dynamics



L1 track trigger for CMS



2010

2012

2014

2016

Parallel Programming Framework and High-speed tomography

Iterative Reconstruction
Laminography
Hough Transform

Discrete Tomography
Clustering

DirectGMA

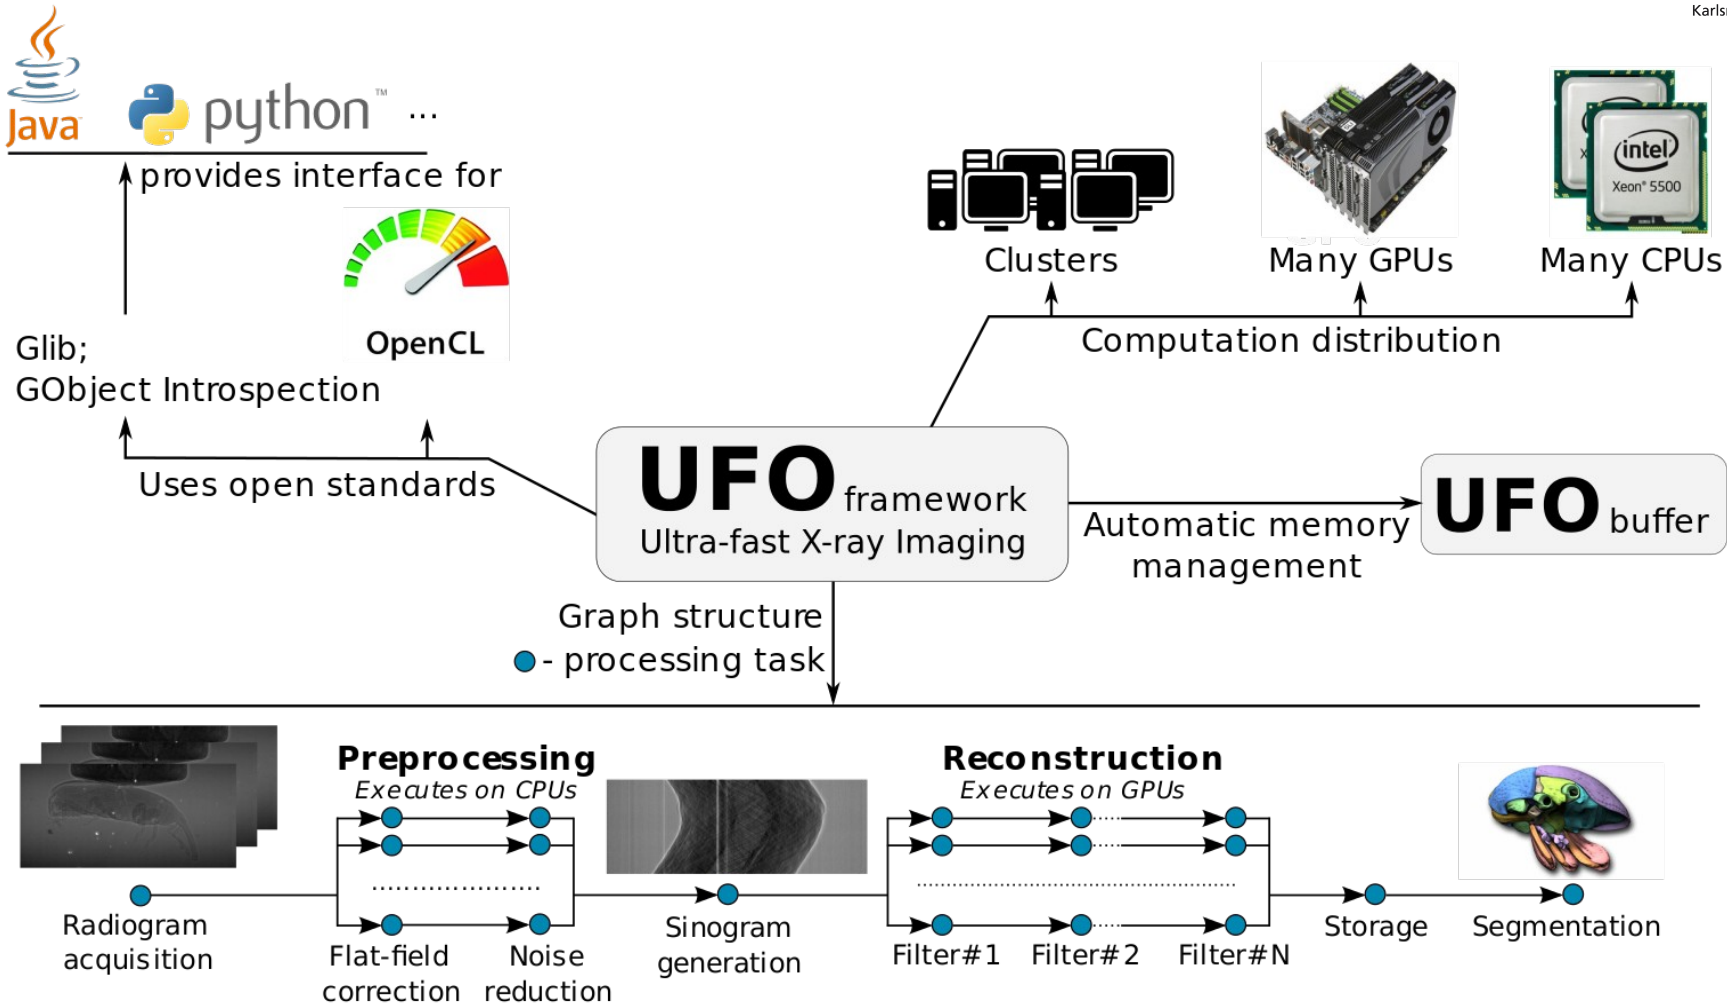
Camera Abstraction Layer
Smart Camera

FPGA Drivers Ecosystem
Virtex7-based DAQ Platform

Concert Control System

Hardware Platform
RDMA Support

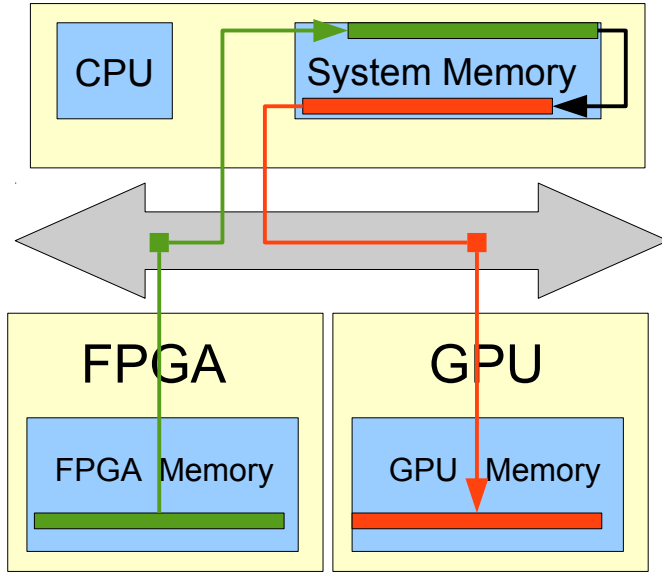
UFO Image Processing Framework



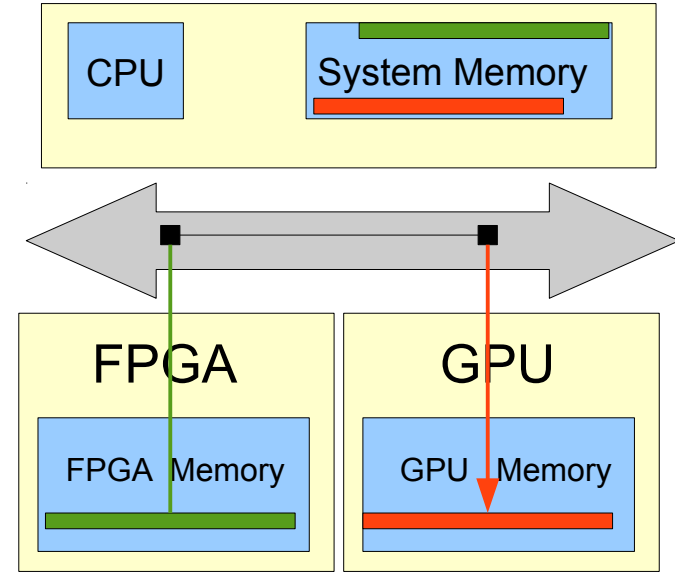
Fully pipelined architecture supporting diversity of the hardware platforms and based on open standards for easy algorithms exchange. Easy prototyping with Python and other scripting languages.

Detector to GPU Transfer

+ 4 memory accesses

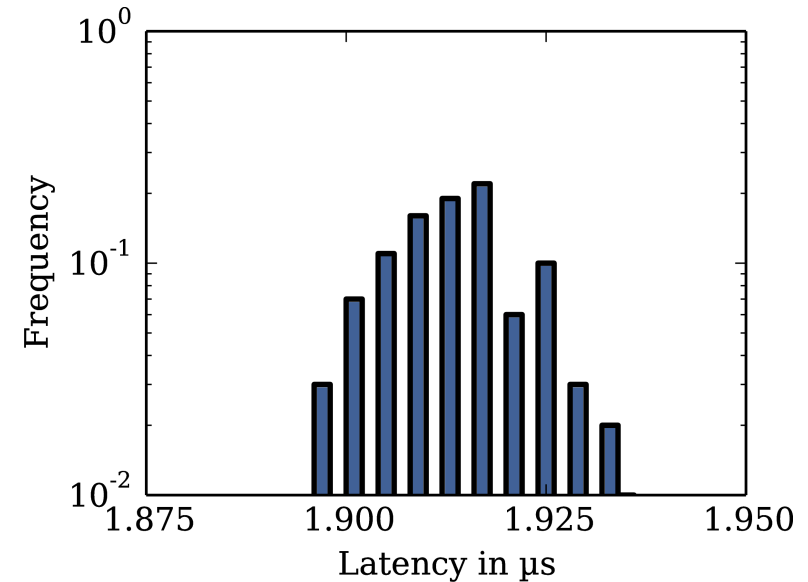



PCIe bus



Standard

GPUDirect



Kernel Launch Penalties

1

CPU-managed

Wait for completion of DMA transfer and launch the CUDA / OpenCL kernel

~25-35 us

2

GPU-managed

Manage DMA in the single-threaded GPU kernel and launch another kernel using Dynamic Parallelism

~25-35 us

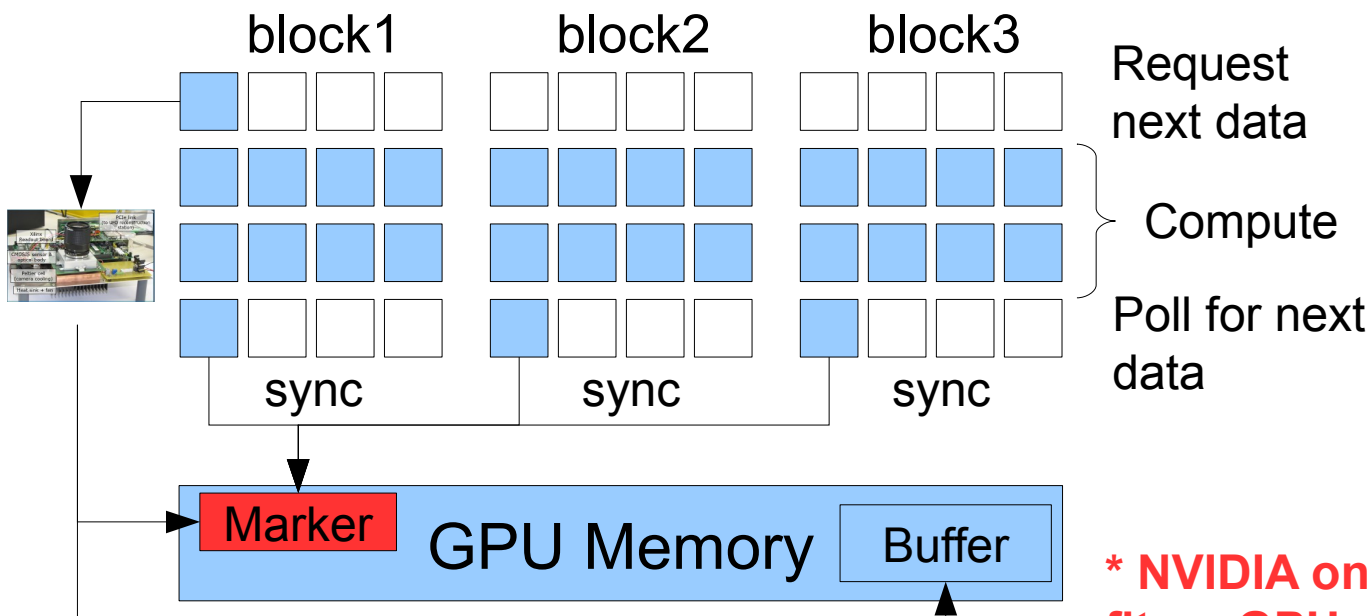
3

Event-based

Block CUDA / OpenCL queue with cuStreamWait32 / clEnqueueWaitSignalAMD until marker is written by DMA engine

~ 8-18 / 25 – 100 us

4 Integrated GPU-side DMA Kernel




(NVIDIA)

(AMD)

Latency vs. blocks

(256 threads per block)

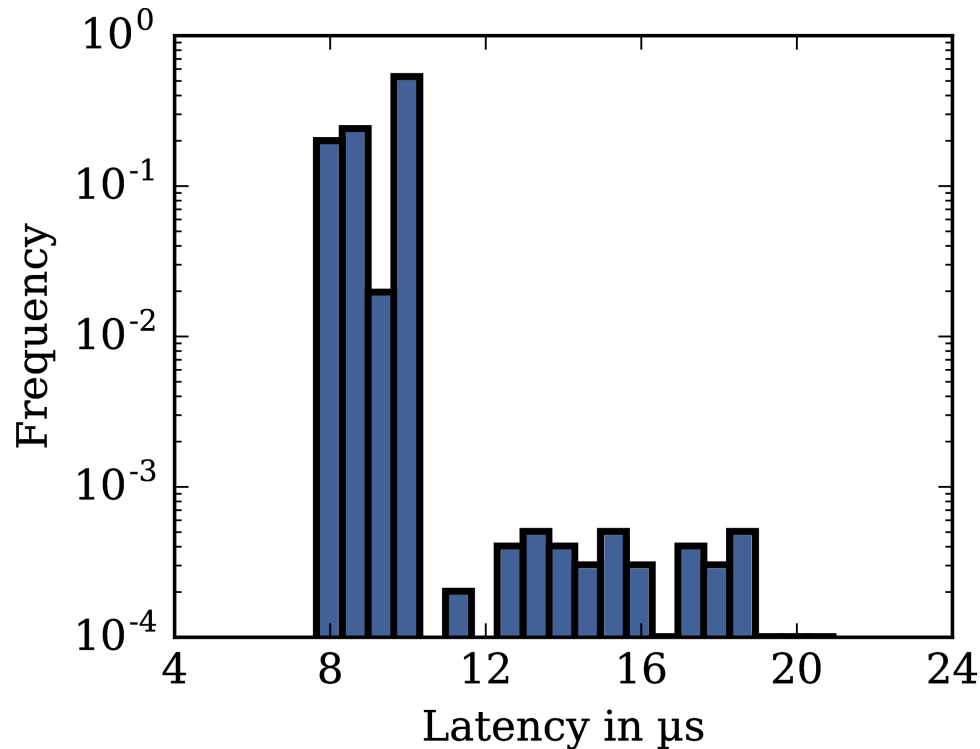


* NVIDIA only and will fail if grid does not fit on GPU completely


The actual latencies including the kernel launch penalties

UFO Framework

Event-based kernel launch




CUDA DMA Kernel

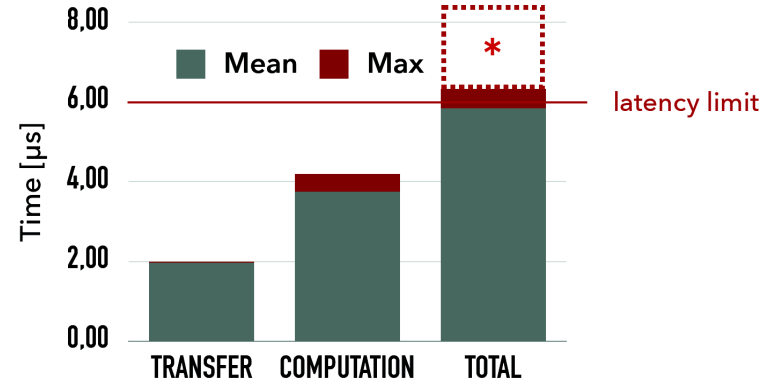


NVIDIA Tesla K40

Prototype of CMS L1 Trigger

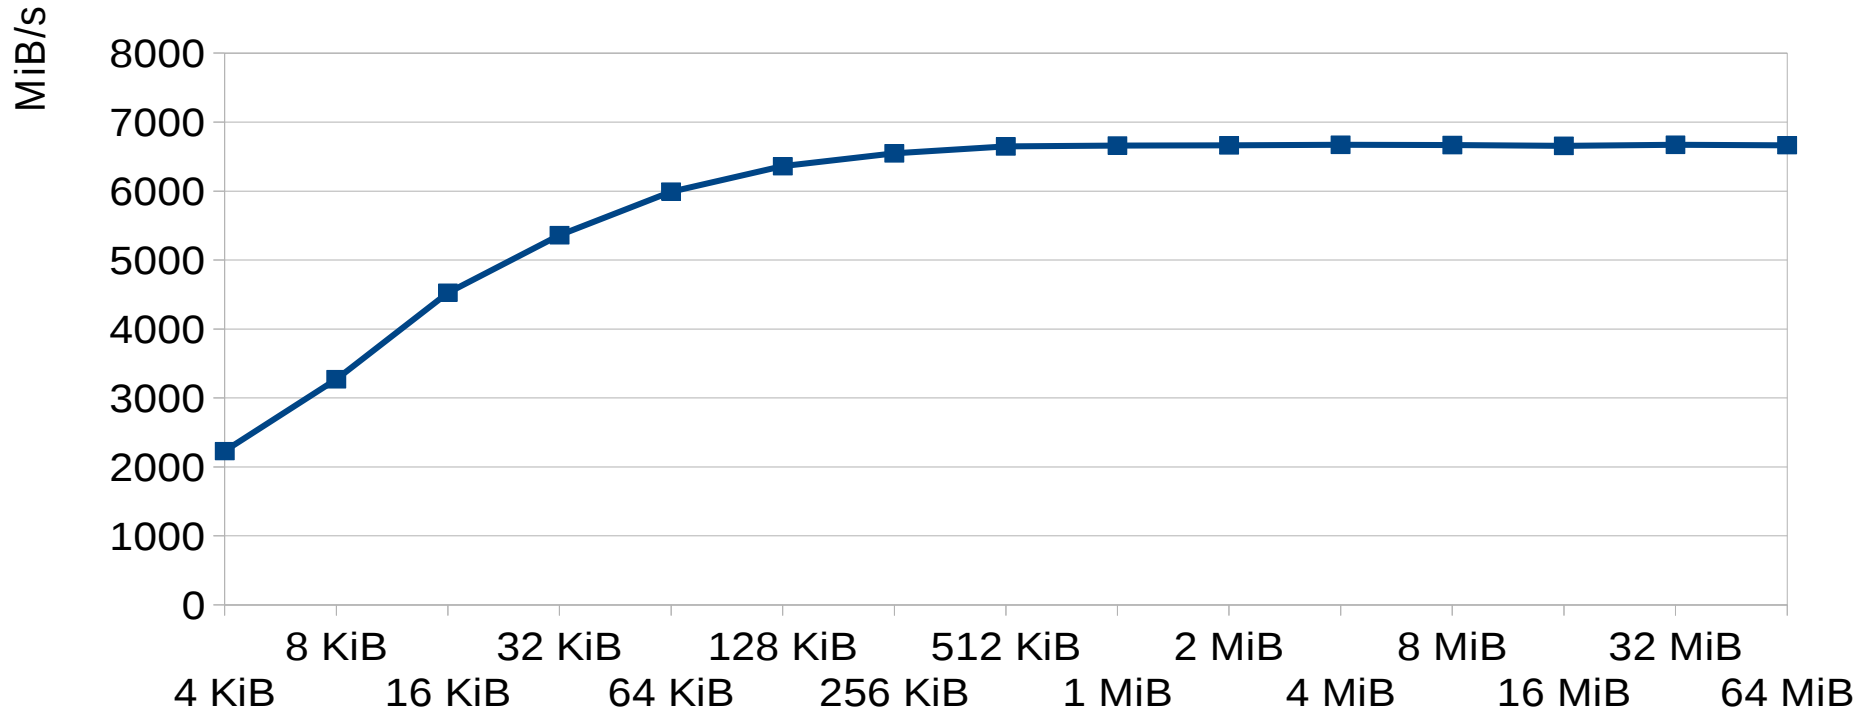


► Read/Uncompress data ► Compute ► Poll



160 stubs, 1 sector

* estimated response time



High-speed streaming to the storage

Network FS

NFS
Samba
SSHFS

Slow


Cluster FS

Lustre (patched kernel)
Gluster
BeeGFS (close-sourced)

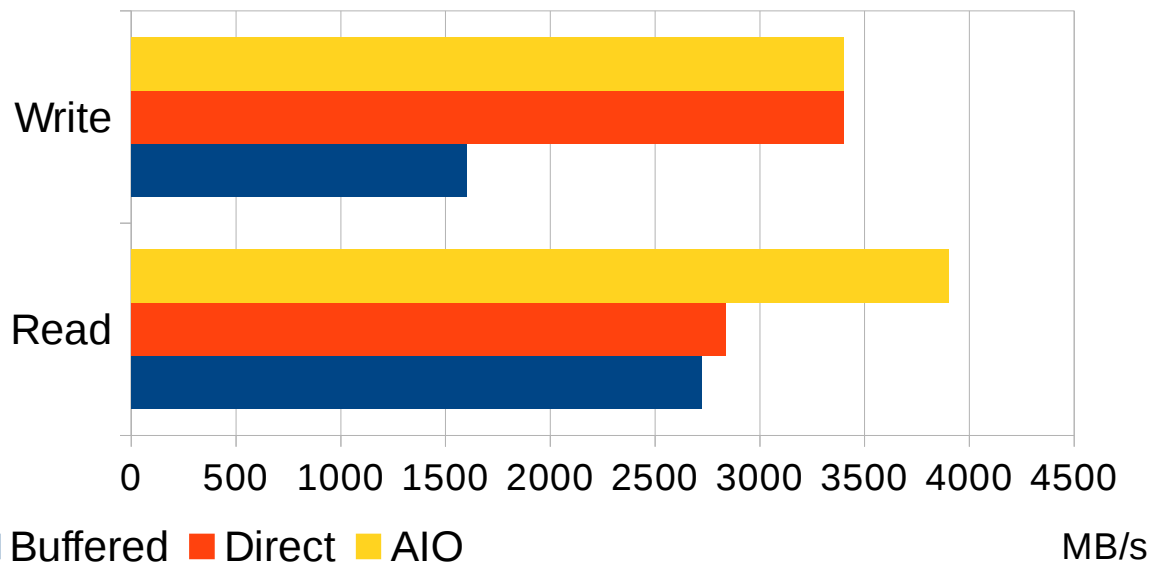
Slow if few nodes

Network Devices

iSCSI (slow)
iSER
OCFS2 (slow)

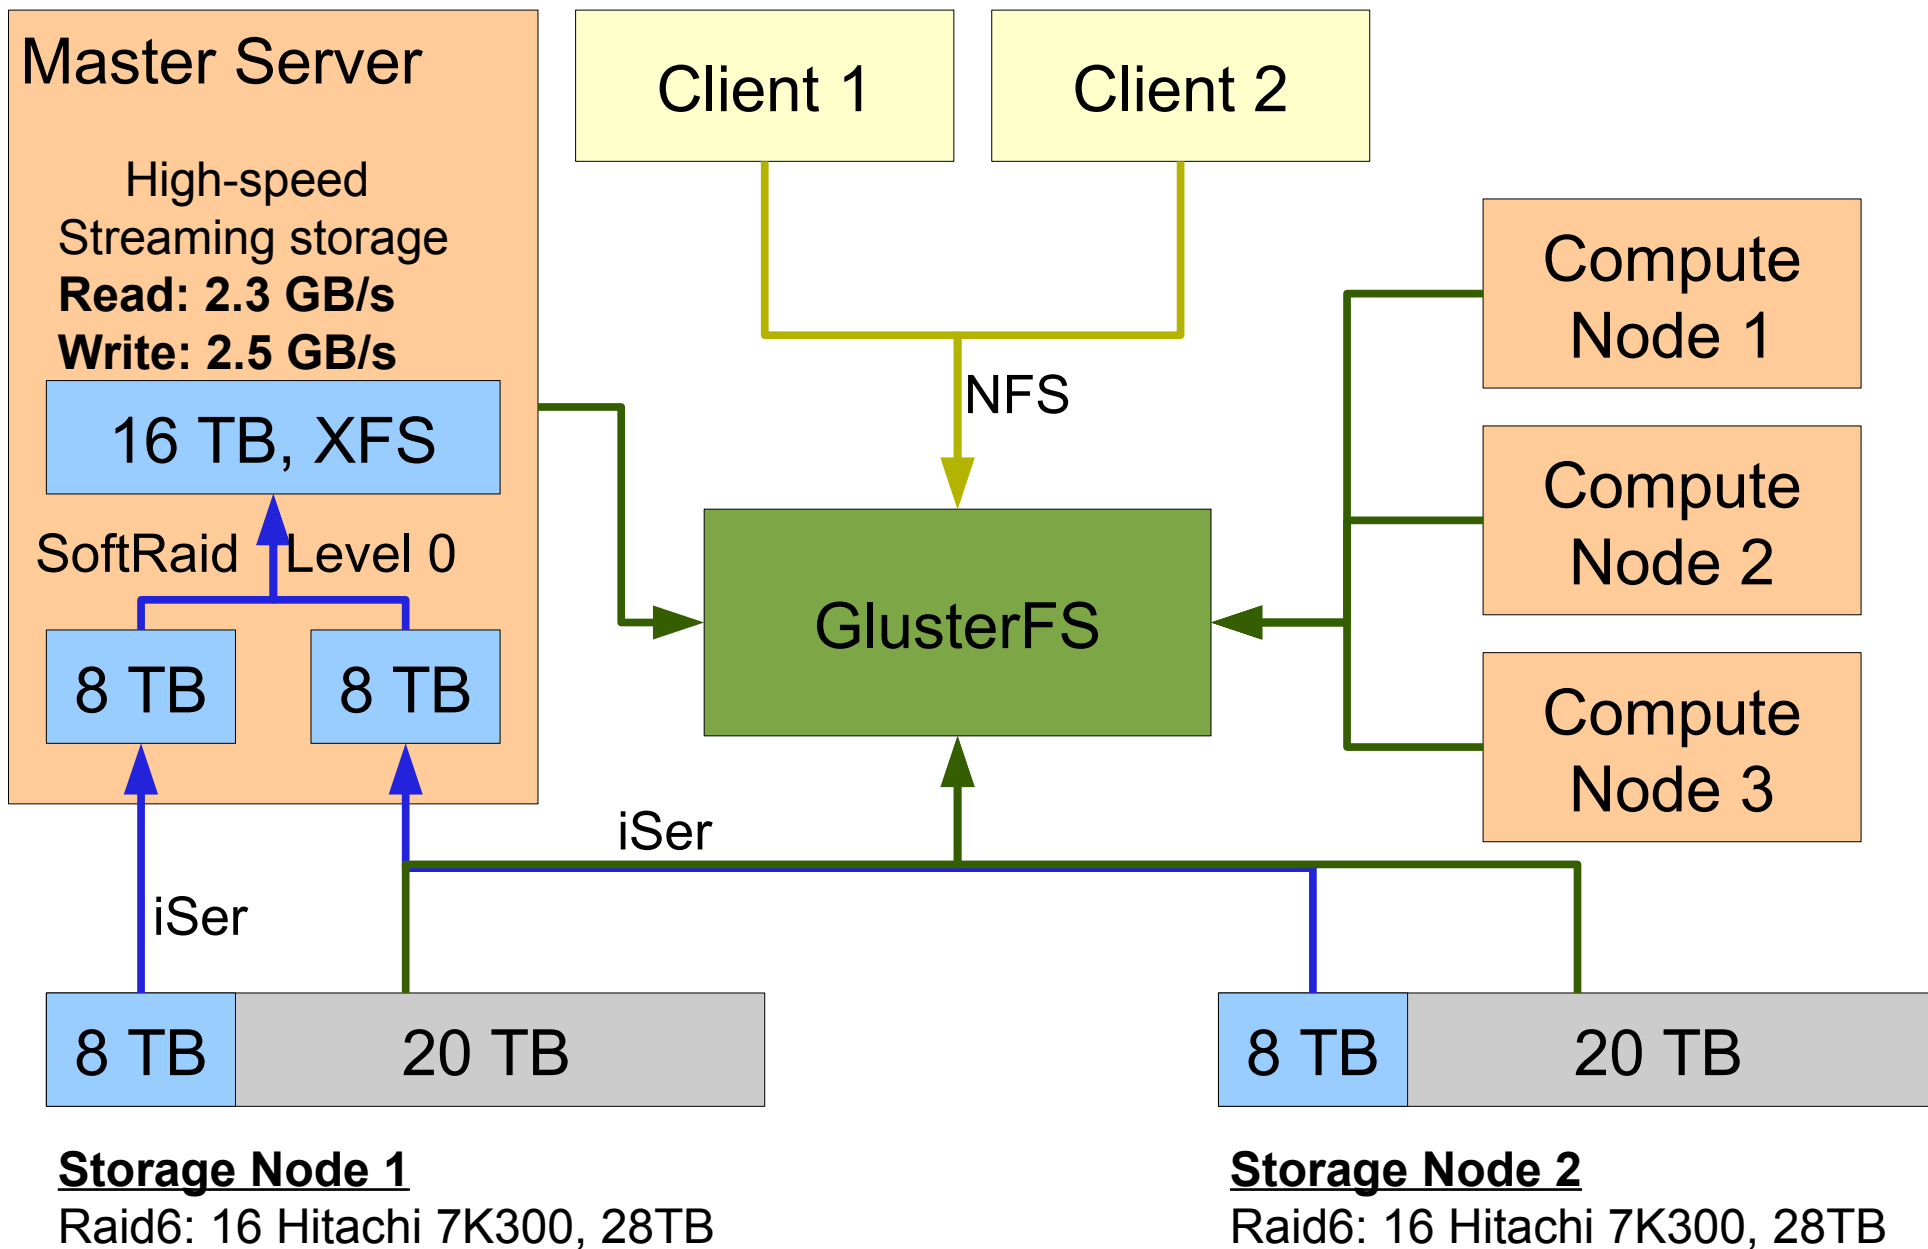


Default data flow in Linux



Optimizing I/O for maximum streaming performance using a single data source/receiver

UFO Storage Subsystem

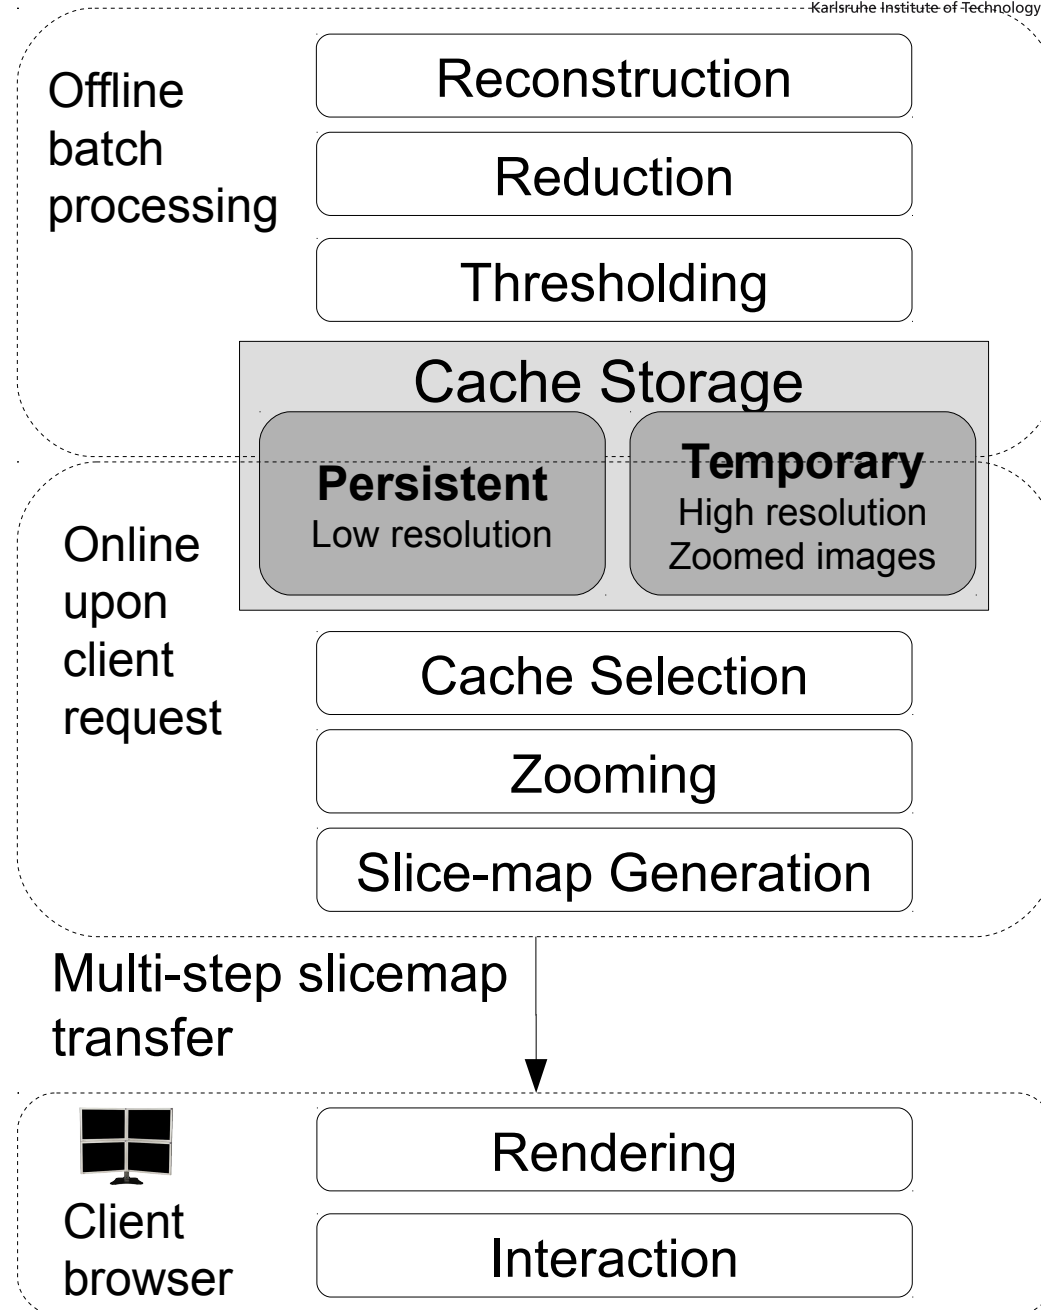


WAVE: Web-based volume visualization

volume visualization



- ▶ Balance between offline jobs, online jobs, and client-side rendering
- ▶ Multi-step transfer to allow quick preview and improve resolution later
- ▶ Multiple zooming levels for inspecting fine details
- ▶ High-quality cuts
- ▶ Multi-modality rendering support



Summary: UFO Software Stack

