Utilization of modern semiconductors ~GPU, SSD, NVRAM, FPGA, ...~

NEC Business Creation Division
The PG-Strom Project
KaiGai Kohei <kaigai@ak.jp.nec.com>

Target of today's discussion

HW Class for Computing

- GPU
- FPGA

HW Class for Storage

- SSD
- NVRAM

Characteristics of computing HW

CPU

- Functional cores, but relatively small number (~20)
- Capability of operating system, storage and network
- Relatively large memory (more than 100GB is usual)

GPU

- Simple cores, but relatively large number (~3000)
- Advantaged on massive numerical operations.
- Programmable, and short time to build and reload (~2sec)
- Relatively small memory (~12GB; GTX TITAN X)

FPGA

- Flexible logic defined by HDL
- Advantaged on known, specific and pre-defined function?
- Programmable, but takes long time to rewrite (~30min)

Characteristics of storage HW

Magnetic Drives

- You know well

SSD

- 200K-400K IOPS, 2.4GB/s throughput
- Interface via filesystem (NVMe also)
- Widely accepted in the market
- No penalty on random access

NVRAM

- RAM speed access
- Interface via memory map
- Not yet commodity in x86_64 server
- No penalty on random access

Interesting Technologies (1/2)

ioMemory Direct I/O API

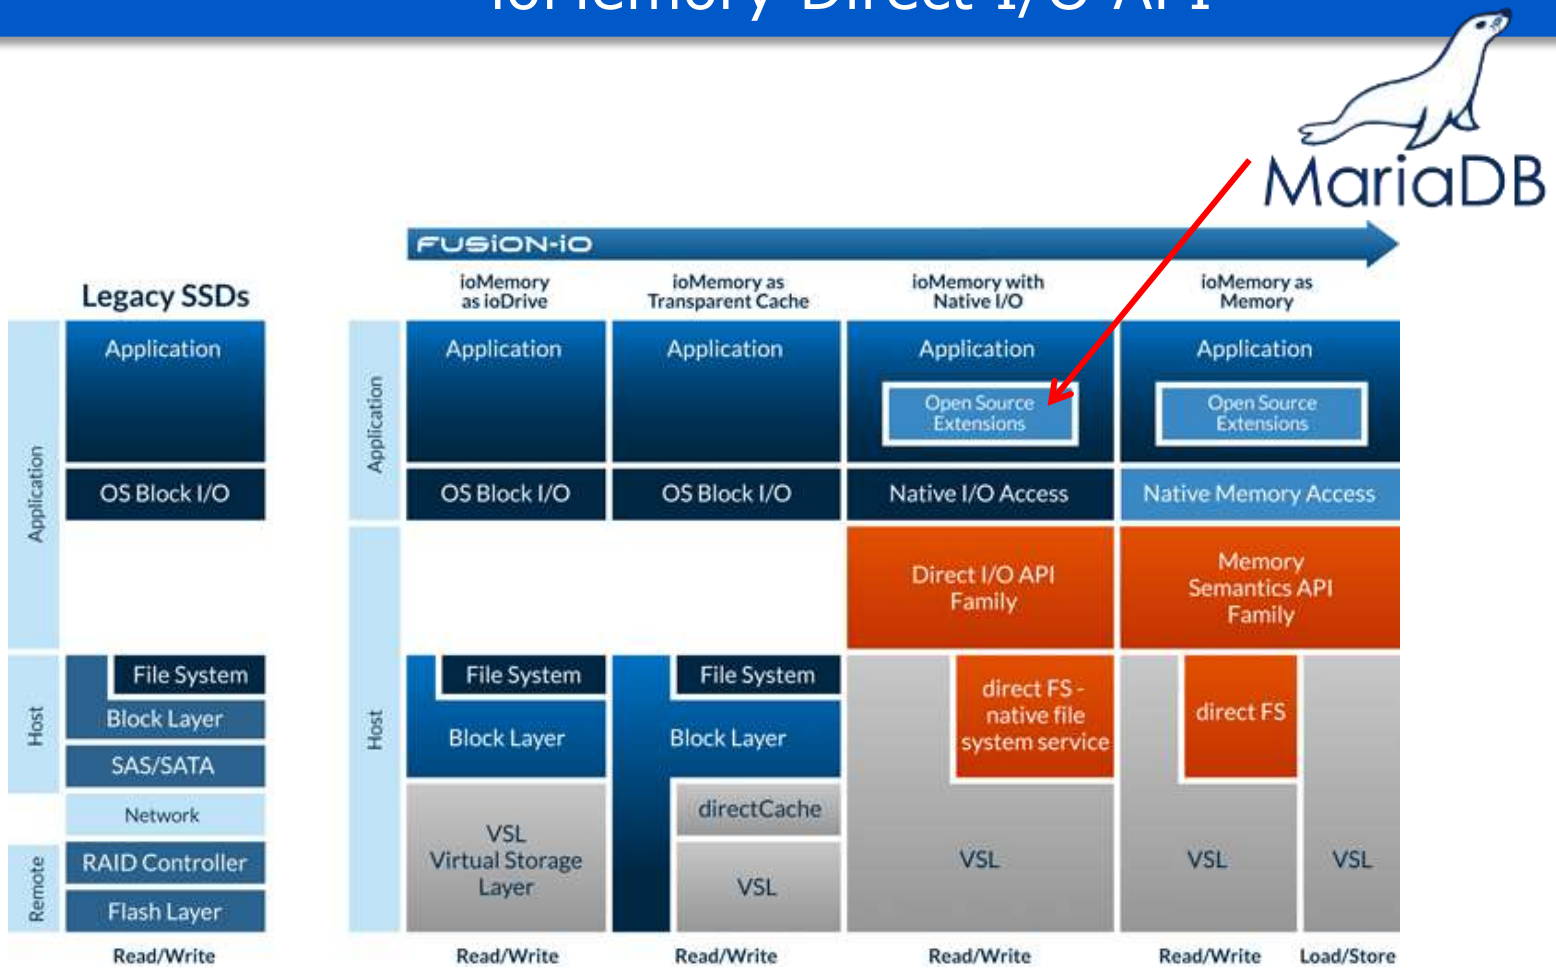

- http://opennvm.github.io/nvm-primitives-documents/
- What is the project status?

Interesting Technologies (2/2)

ioMemory (Fusion-IO) <-> GPU Direct Memory Access

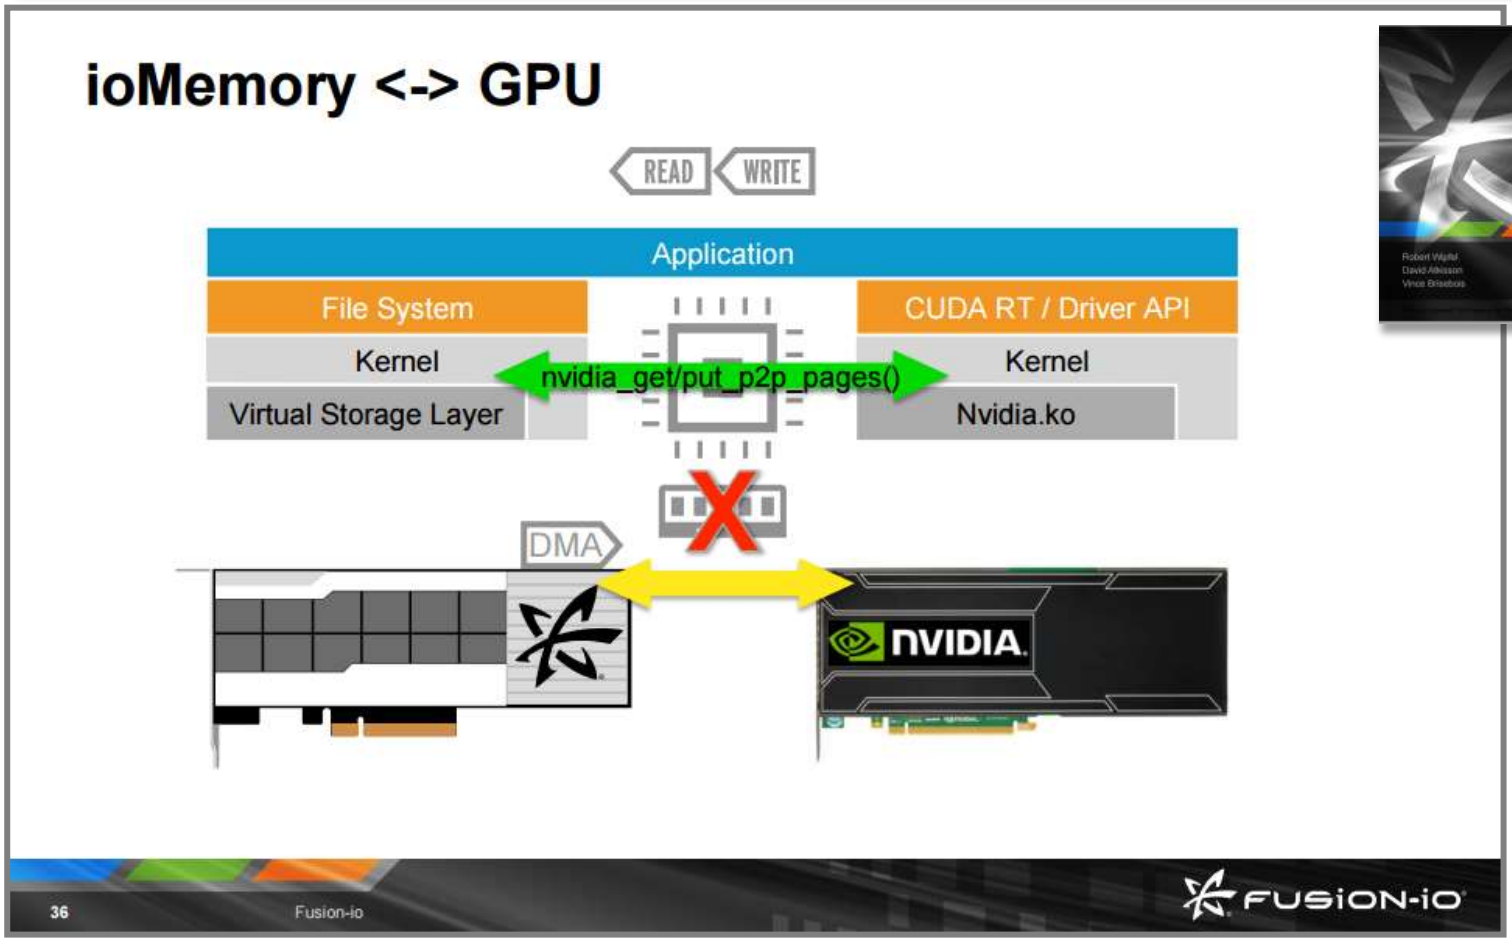

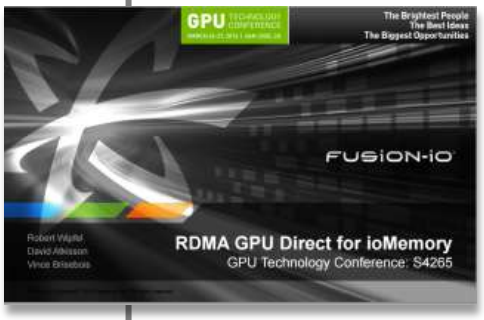

SOURCE: RDMA GPU Direct for ioMemory, Robert Wipfel, David Atkisson, Vince Brisebois
GPU Technology Conference: S4265

Ideas towards PostgreSQL adoption

HW class for computing

- Off-loads of CPU intensive workloads
  - ➔ Join, Sort, Aggregate on CustomScan node?
- Procedural Language support
  - ➔ pl/CUDA, pl/FPGA?

HW class for storage

- No random access penalty
  - ➔ suitable parallel scan, but planner may needs to pay attention
- Higher read throughput for OLAP workloads
  - ➔ I/O density of single query execution is concern
- Small latency for transaction logs
  - ➔ NVRAM on transaction log buffer, or atomic write to SSD

Proprietary Tools

- CUDA, ioMemory SDK, vendor specific drivers, ...etc

My Vision (1/2)

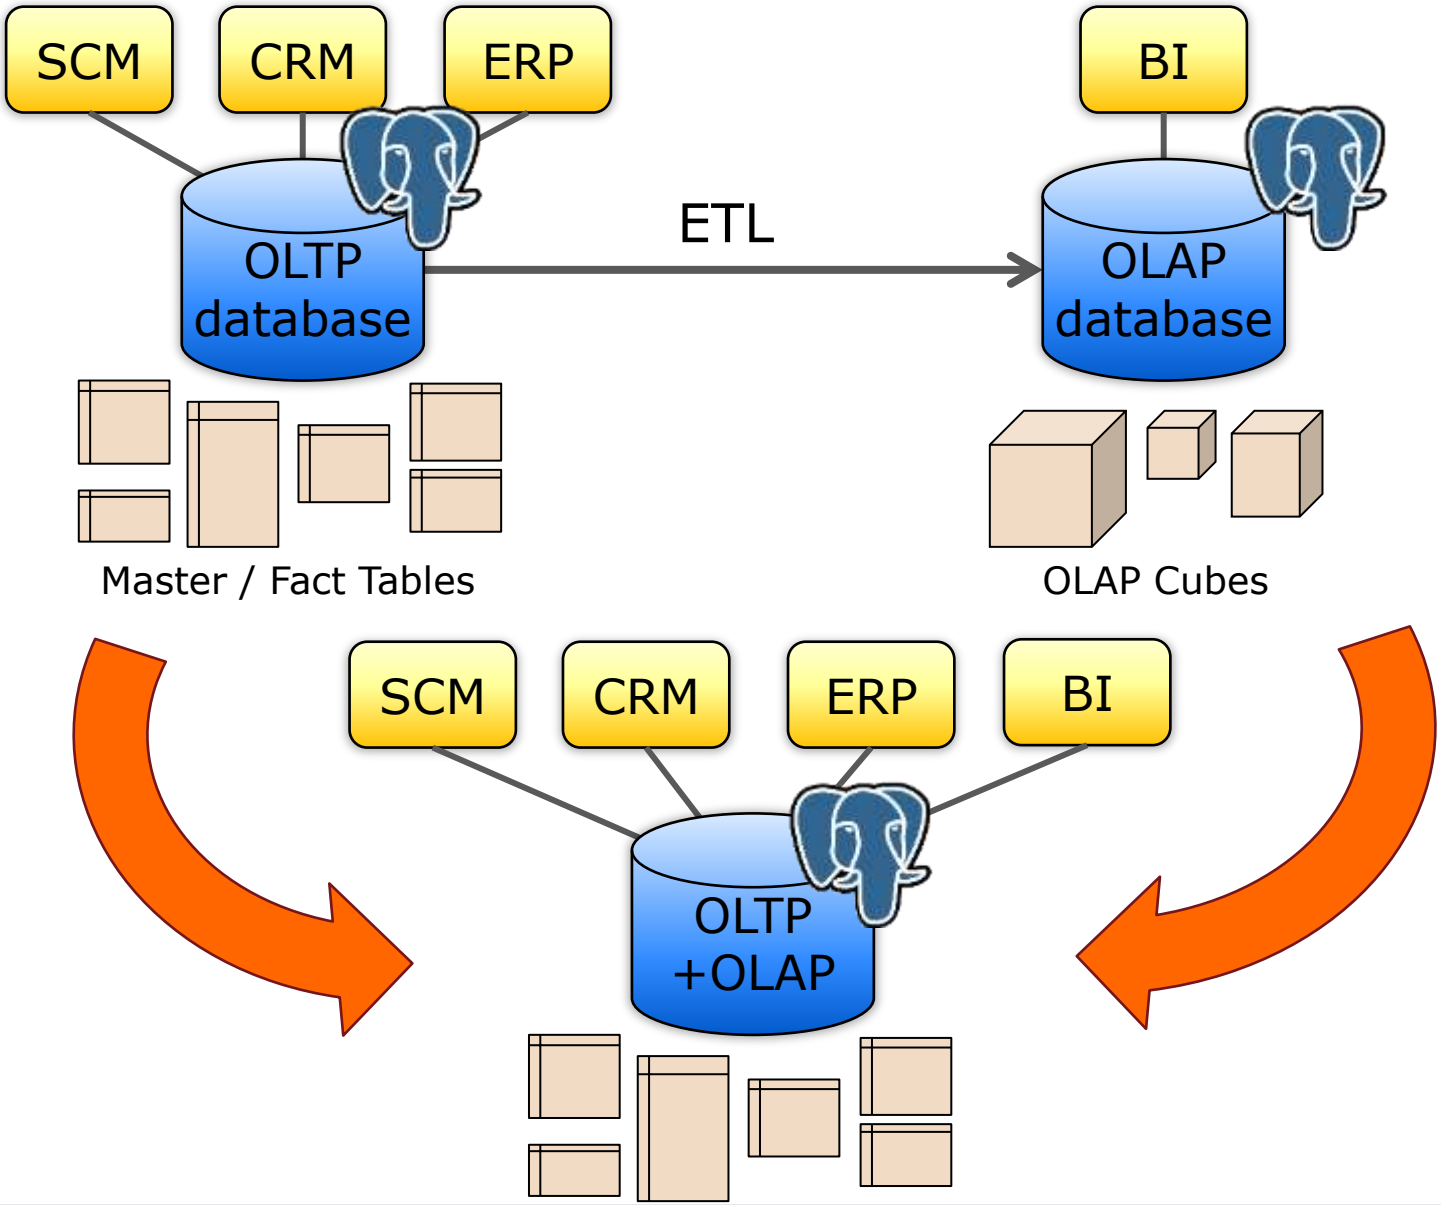

My Vision (2/2)

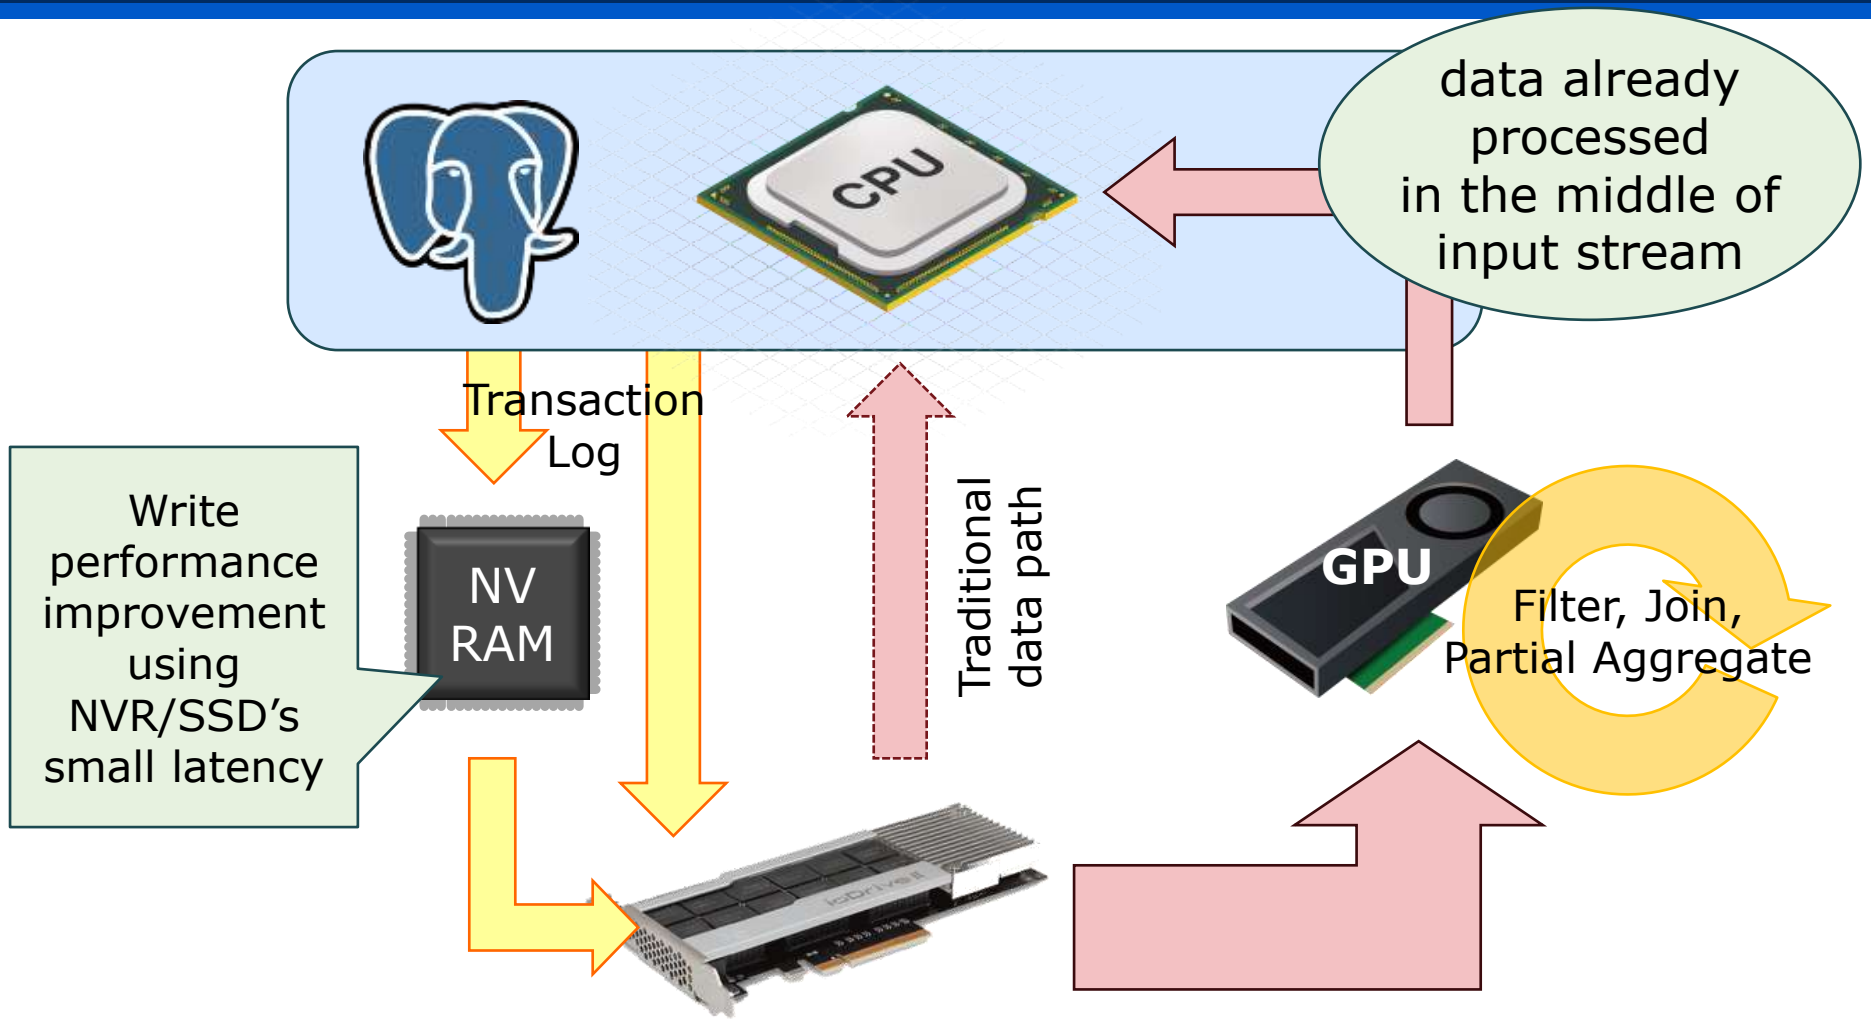

- Cost: Server=$20K, SSD=$10K, GPU=$5K, NVRAM=$5K?

➔Total: $40K + "value of PostgreSQL"

\Orchestrating a brighter world

NEC

NEC brings together and integrates technology and expertise to create
the ICT-enabled society of tomorrow.
We collaborate closely with partners and customers around the world,
orchestrating each project to ensure all its parts are fine-turned to local needs.

Every day, our innovative solutions for society contribute to
greater safety, security, efficiency and equality,
and enable people to live brighter lives.