# DDN | SFA10KE™

## Embedded processing with iRODS and SFA10KE

J.J. Krol,
Systems Engineer

DataDirect™ NETWORKS

# Who is DataDirect Networks?

## Key Statistics

- Delivers highly scalable and highly efficient storage solutions that enable customers to accelerate time to results, scale simply as data sets continue to grow, and gain competitive advantage through resolving performance and capability scaling challenges
- **Established:** 1998
- **Financials:** Over $200M Annually, Profitable and Growing
- **Headquarters:** Chatsworth, California USA
- **Employees**: Approximately 400 Worldwide
- **Customers:** Over 1,000 Worldwide
- **Footprint:** 17 Industries, 4 Continents, 49 Countries
- **Go to Market:** Global Partners, VARs, Resellers
- **Key Market Segments:**
  - High Performance Computing & Life Science
  - Cloud & Web Content
  - Rich Media
  - Intelligence/Federal
  - Surveillance

## Industry Validation

| | |
|---|---|
| IDC Analyze the Future | World's Largest Privately-Held Storage Co. |
| Deloitte. | Fast500 Technology Company |
| Inc. | Inc. 500\|5000 High-Growth Company |
| HPCwire | Best HPC Storage Platforms |
| Frost & Sullivan | Best Practice for Digital Media |

## 1000+ World-Leading Customers

SONY · AIRBUS · Microsoft · Oak Ridge National Laboratory · yousendit · Warner Bros. · General Dynamics · IBM · Total · photobox · Boeing · Mediaset



DataDirect NETWORKS

# What is ‘Embedded Processing’?

**And why ?**

- **Do data intensive processing as ‘close’ to the storage as possible.**
  - Bring computing to the data instead of bring data to computing
- **HADOOP is an example of this approach.**

- **Why Embedded Processing?**
- **Moving data is a lot of work**
- **A lot of infrastructure needed**

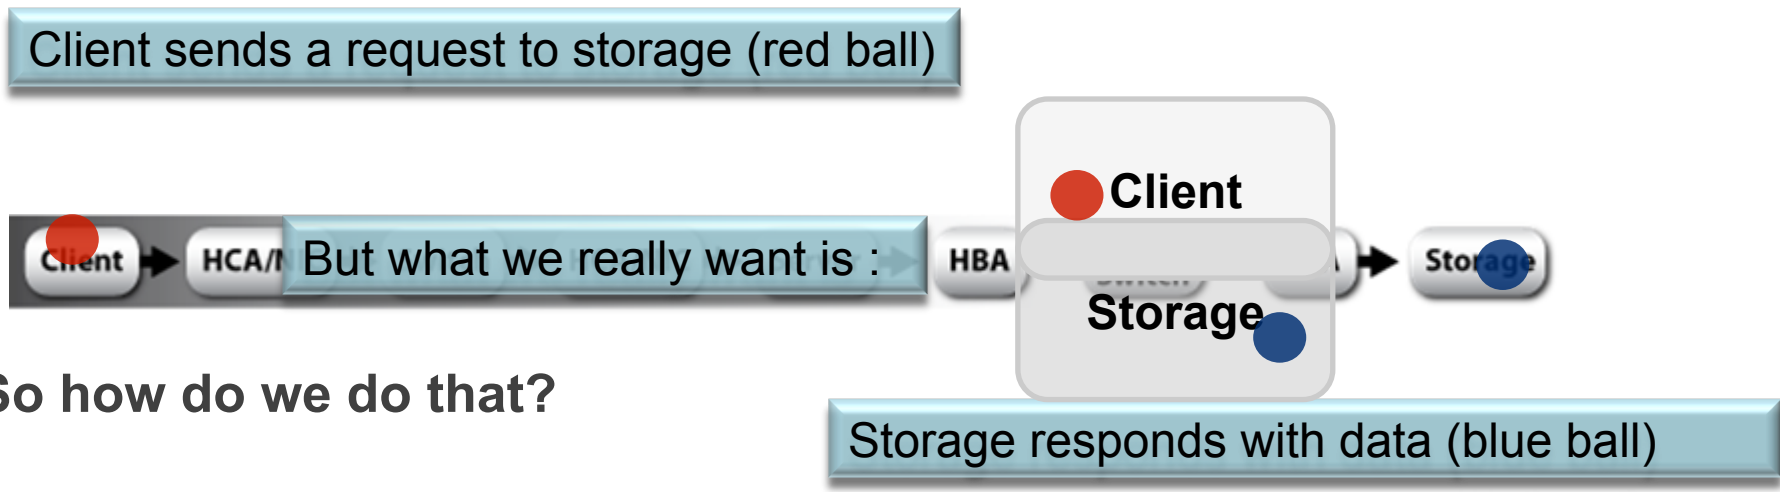


- **So how do we do that?**



DataDirect NETWORKS

# Storage Fusion Architecture (SFA)

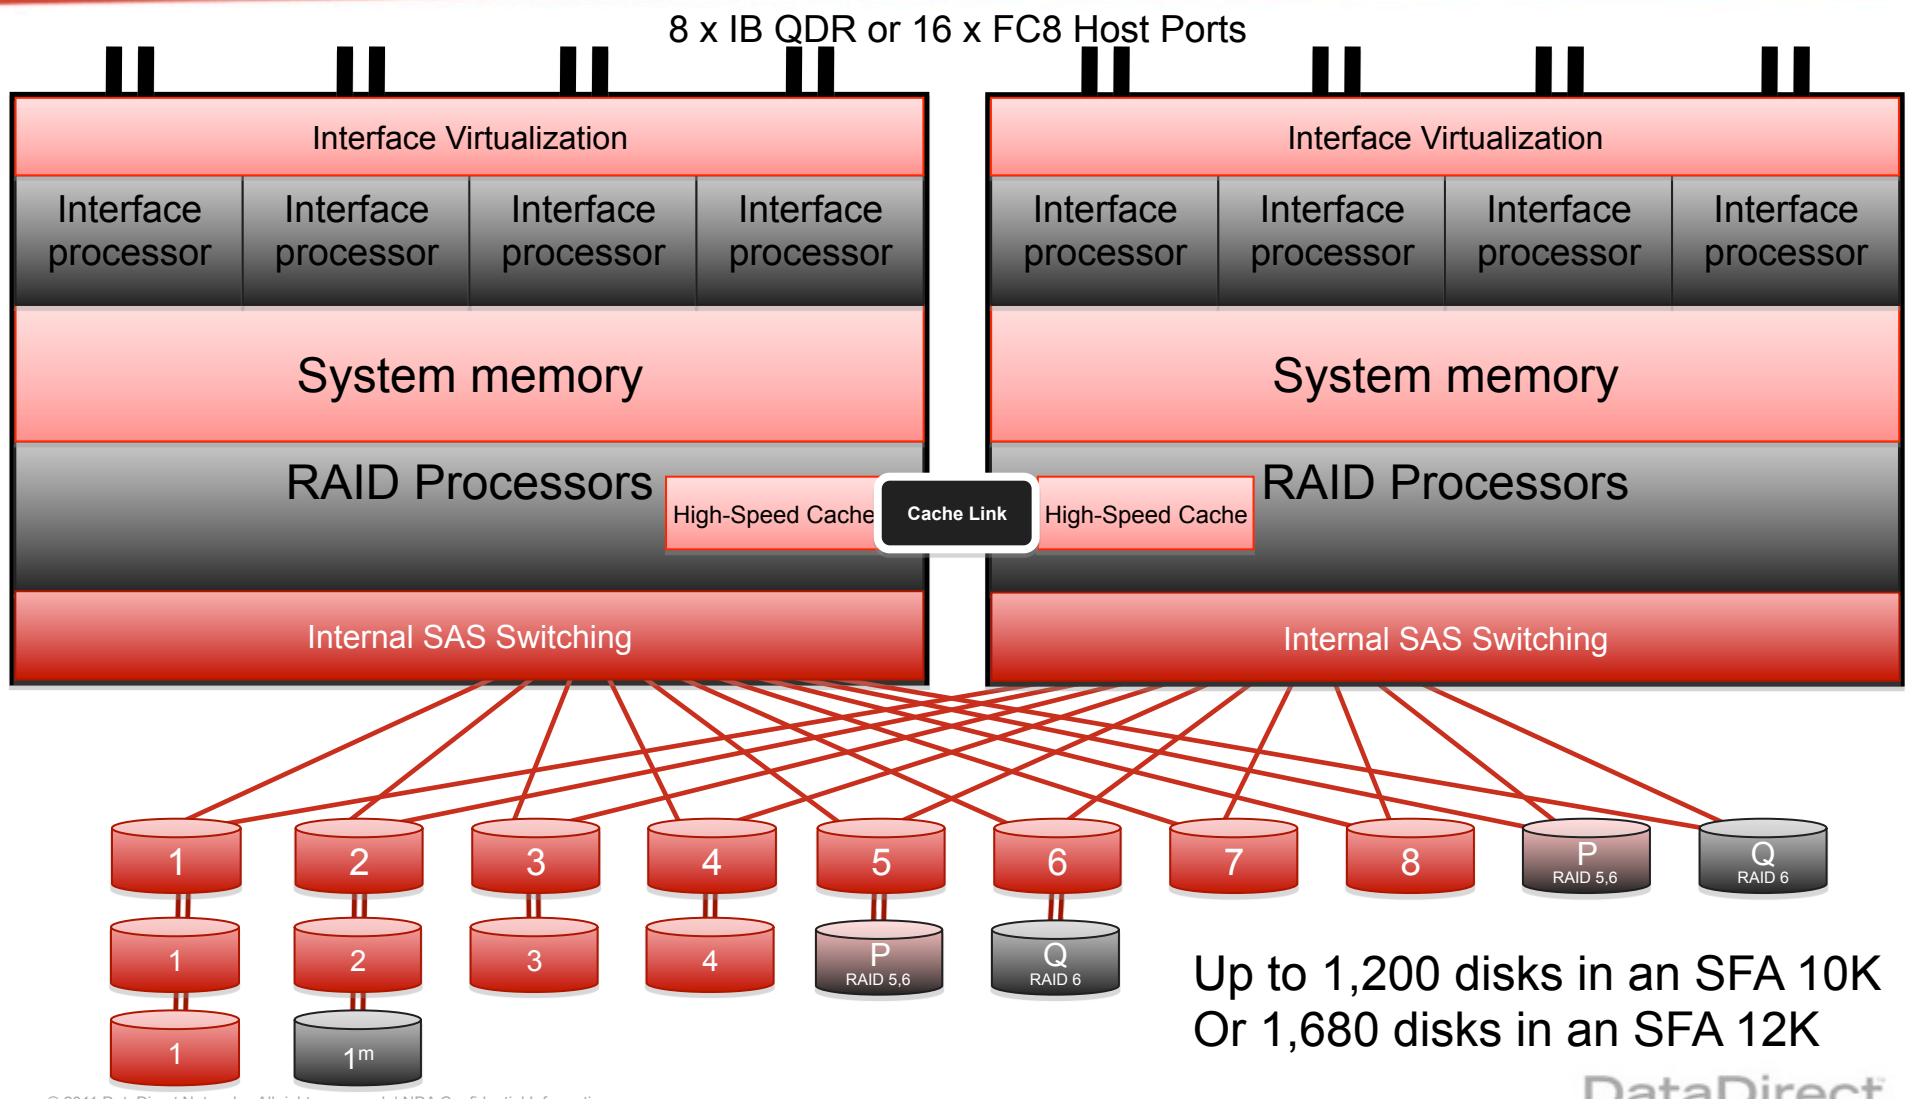


Up to 1,200 disks in an SFA 10K
Or 1,680 disks in an SFA 12K



DataDirect NETWORKS

# Repurposing Interface Processors

- In the block based SFA10K platform, the IF processors are responsible for mapping Virtual Disks to LUNs on FC or IB

- In the SFA10K**E** platform the IF processors are running VMs
- The OS running on those VMs uses a driver to access the RAID processors directly
- RAID processors place data (or use data) directly in the VM's memory
- One hop from disk to VM's memory

- Now the storage is no longer a block device
- It is a storage appliance with processing capabilities

# One SFA-10KE controller

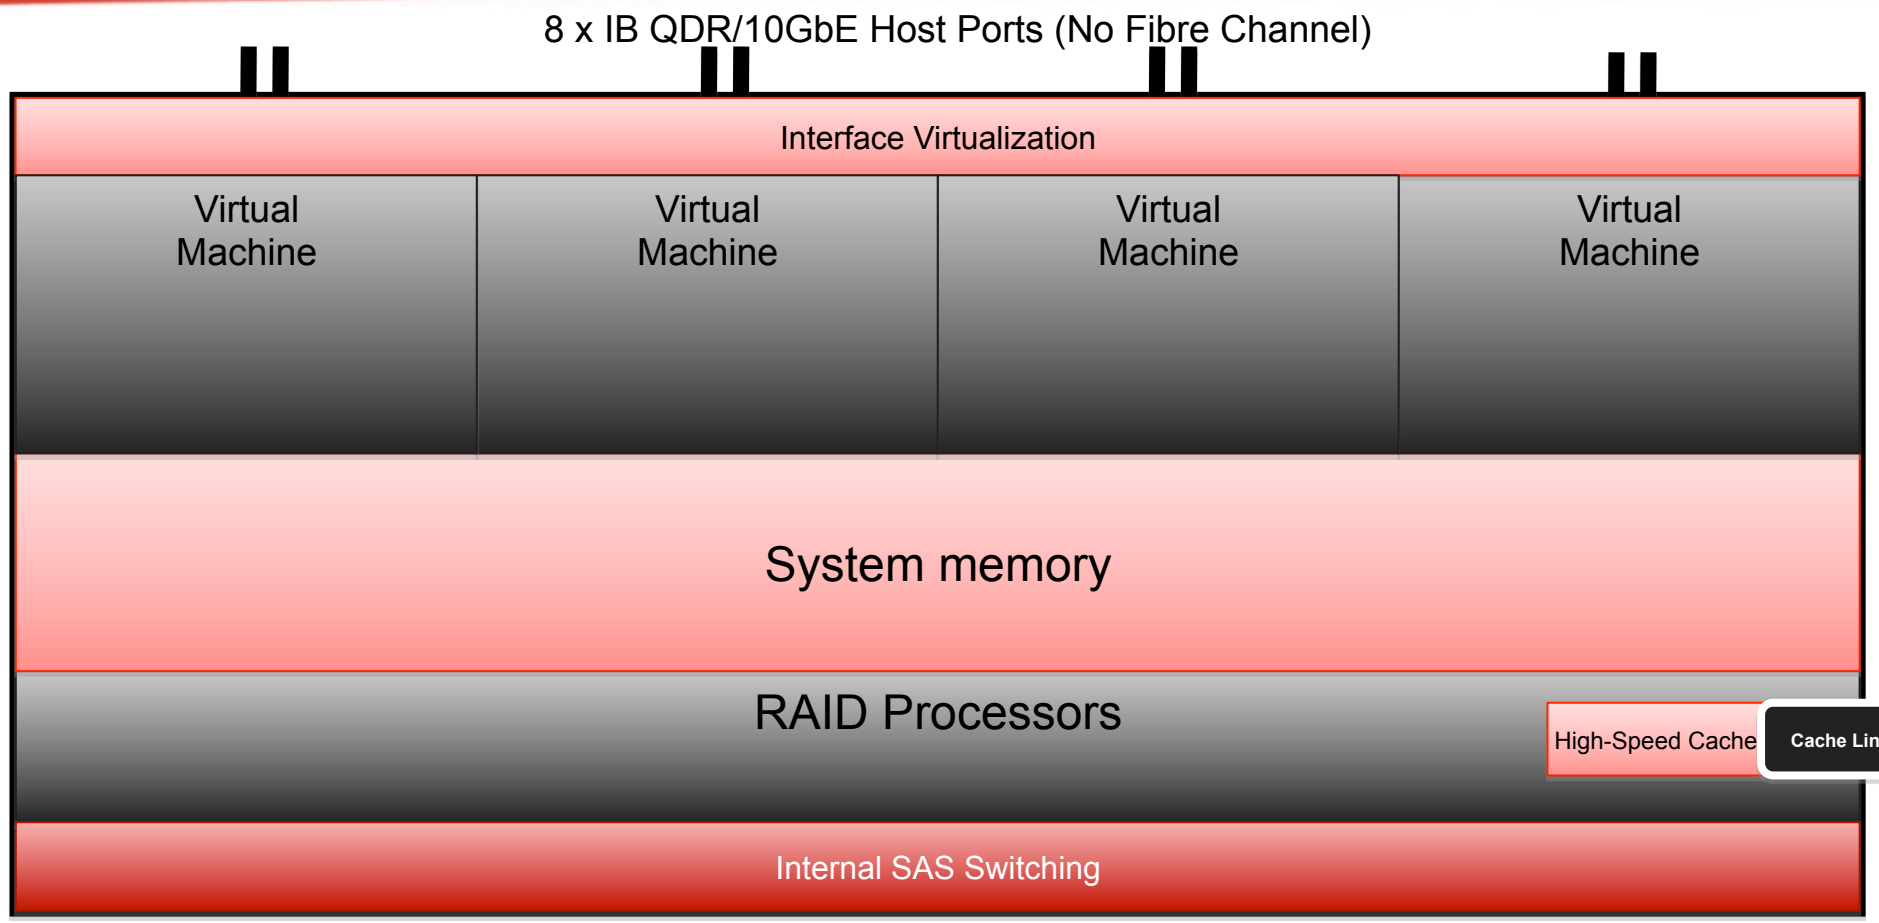




DataDirect NETWORKS

# Example configuration

- **Now we can put iRODS inside the RAID controllers**
  - This give iRODS the fastest access to the storage because it doesn't have to go onto the network to access a fileserver. It lives **inside** the fileserver.
- **We can put the iCAT on a separate VM with lots of memoty and SSDs for DB storage**
- **Either use all VMs for iRODS or add a parallel filesystem such as GPFS for fast scratch**
- **The filesystem uses SAS for frequent used files and SATA for the rest**
- **The following example is a mix of iRODS with GPFS**
  - The same filesystem is also visible from an external compute cluster via GPFS running on the remaining VMs
- **This is only one controller, the 4 VMs on the other controller need some work too**
  - They see the same storage and can access it at the same speed.
- **On the SFA-12K we will have 16 VM's available running on Intel Sandy Bridge processors.** **(available Q3 2012)**

# Example configuration

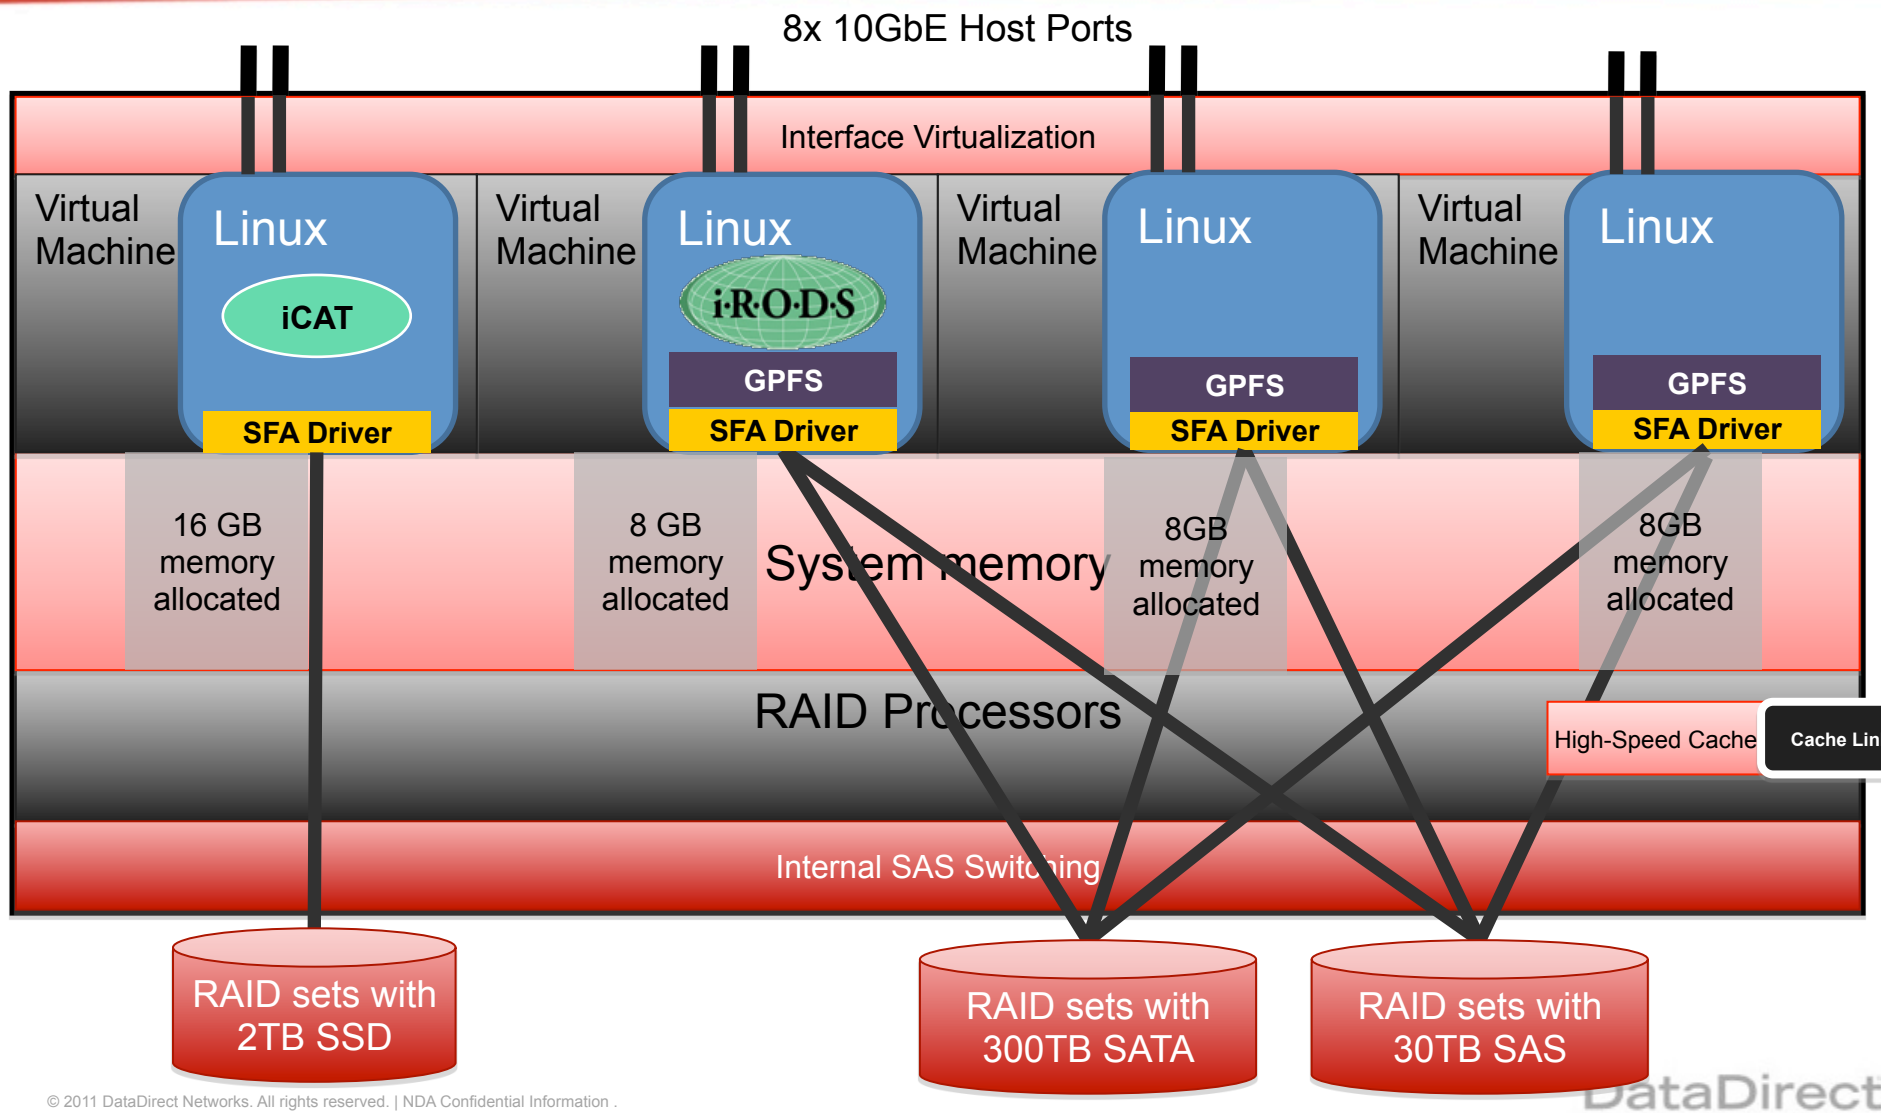




DataDirect NETWORKS

# Running Micro Services inside the controller

- Since iRODS runs inside the controller we now can run iRODS MicroServices right on top of the storage.
- The storage has become an iRODS appliance ‘speaking’ iRODS natively.
- We could create ‘hot’ directories that kick off processing depending on the type of incoming data.

# DDN | SFA10KE™
## With iRODS and GridScaler parallel filesystem

An iRODS rule requires that a copy of the data is sent to a remote iRODS server

Reading data via parallel paths with DDN | GRIDScaler embedded in the system.

Writing the result of the processing

iRODS server.

Remote iRODS server

Computing

All of this happened because of a few rules in iRODS that triggered on the incoming data.

In other words, the incoming data drives the processing.

server

DataDirect NETWORKS

After the conversion another MicroService submits a processing job on the cluster to process the uploade

was configured with a rule to

Here is a configuration with iRODS for datamanagement and GridScaler for fast scratch space. Data can come in from clients or devices such as sequencers.

(click to continue)

Virtual Machine 4

Meta data

Registering the da

# Thank you!

## Questions??

**Jan Jitze Krol**
**jkrol@ddn.com**



DataDirect NETWORKS