NFS INDUSTRY CONFERENCE

NFS Clusters

Xiaoye Jiang

EMC Corporation

jiang_xiaoye@emc.com

NFS INDUSTRY CONFERENCE

NFS Clustering – Overview

- Approaches to Scaling NFS
- Enabling Technology: HighRoad
- Bandwidth and Capacity on Demand

NFS INDUSTRY CONFERENCE

Approaches to Scaling NFS

Monolithic Servers

+ Simple Scaling w/ Hardware

- Single Point of Failure

- Practical

Active or Passive Pairs

+ No SPOF

- Double

- No Perf

True Clusters

+ Simple Scaling w/Hardware

+ Bandwidth on Demand

+ Best Price/Performance

NFS INDUSTRY CONFERENCE

Cluster Challenges

Cache Coherency among nodes!

Often the penalty paid in performance and complexity is greater than the benefit received.

i.e. Oracle OPS / RAC

Two approaches to resolve this basic cluster challenge –

- Distributed Lock Managers
- Meta Data Servers

NFS INDUSTRY CONFERENCE

HighRoad – (MPFS)

- HighRoad is the foundation for a true R/W cluster
  - Allows Bandwidth on Demand
  - Consistent Locking
  - User view of ‘one large server’

NFS INDUSTRY CONFERENCE

HighRoad

The Basic Technology

- NFS cluster is based on HighRoad
  - HighRoad originally offered as a client side solution
  - Separates Data from Control
    - Routes Data over the SAN
    - Routes Control over the IP Network

The HighRoad Topology

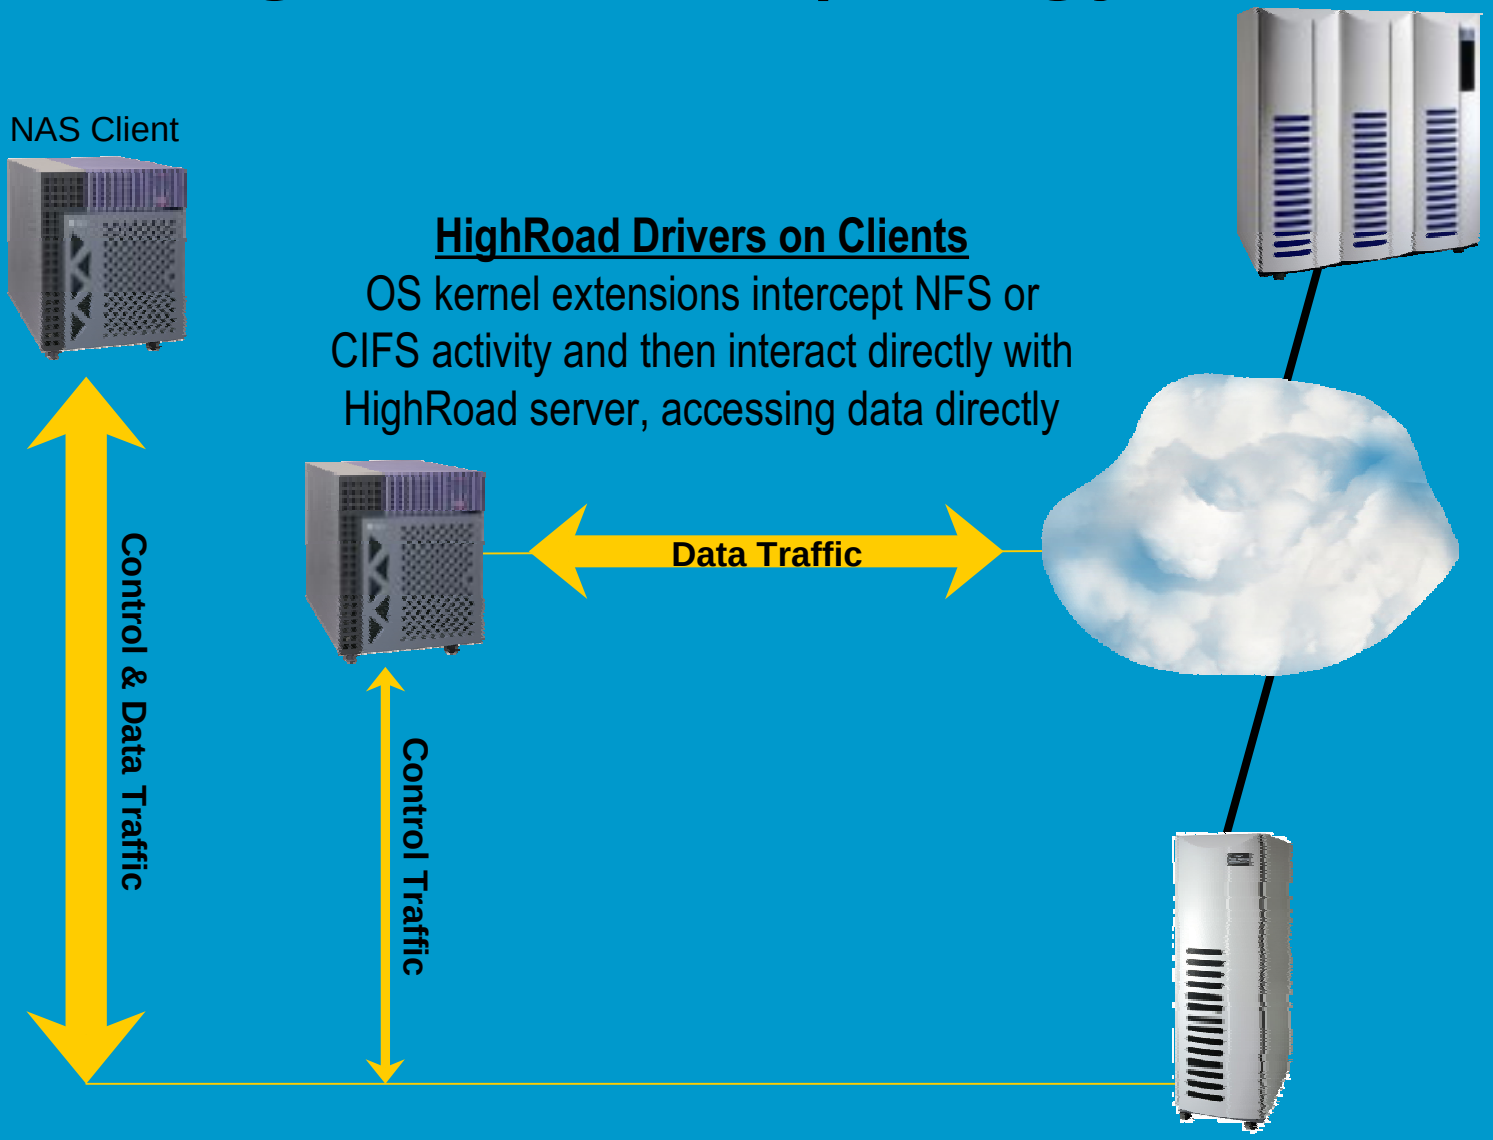

NFS INDUSTRY CONFERENCE

NFS Cluster Concept

- Port the client software to the real time OS that runs on the Data Movers
- Have multiple secondary Data Movers act as clients of a primary Data Mover
- Terminology
  - Primary Node: Data Mover that works as HighRoad Server. One Per Cluster filesystem. Provides Metadata to Secondary Nodes.
  - Secondary Node: Data Mover that works as HighRoad Client. Obtains MetaData from Primary Node. Perform Data I/O Locally.

NFS INDUSTRY CONFERENCE

NFS Cluster Basics

- Primary Node Owns the File System
- Local File Requests
  - Served by Primary Node as normal NFS requests
- Remote File Requests
  - HighRoad Protocol
  - MetaData Obtained from Primary
  - Data Transferred directly from Local Storage
- Transparent to Clients

HighRoad Cluster Topology

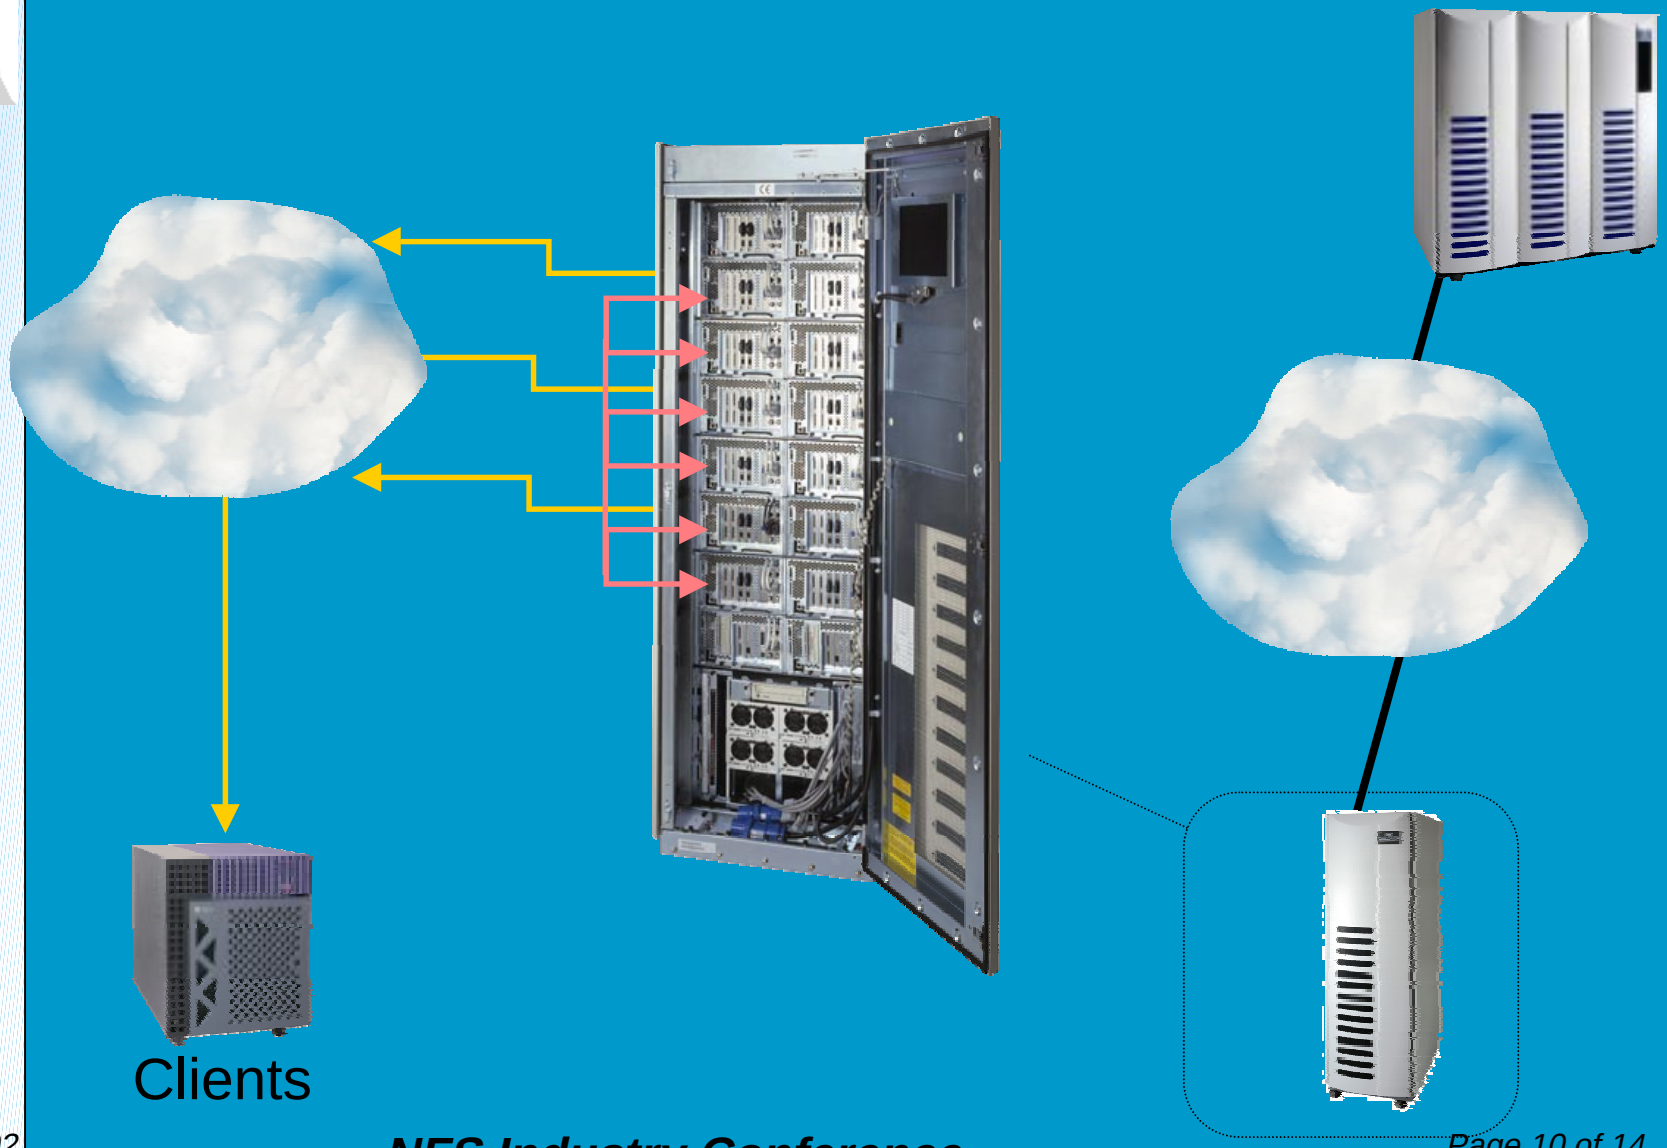

HighRoad Inside the Fileserver

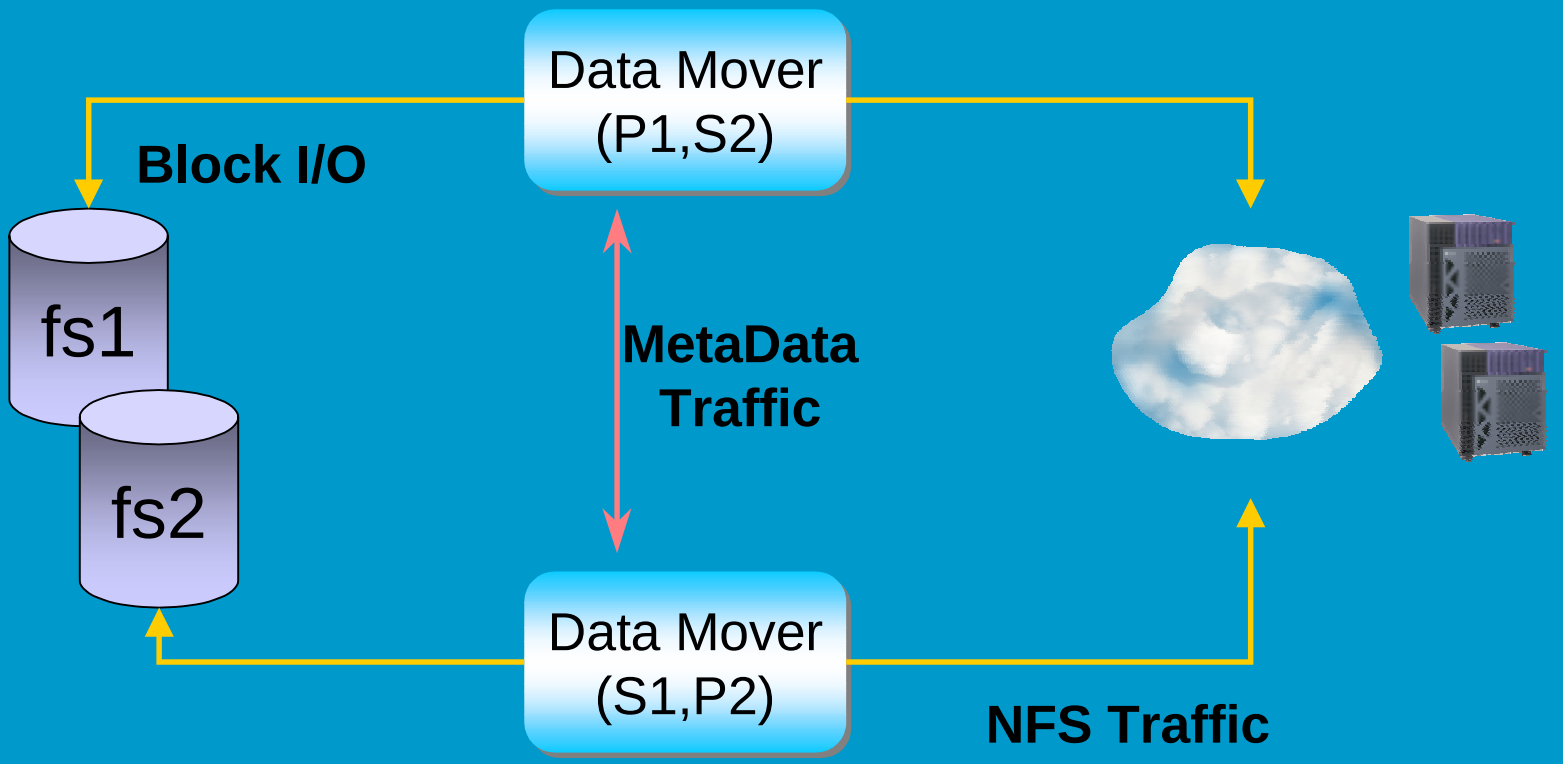

NFS INDUSTRY CONFERENCE

NFS Cluster Advantages

- NFS access to a R/W Shared File System
- Transparent to NFS client
- Data Consistency Maintained by HighRoad Technology
- Lockd Supported
- User/Group Quota Supported
- CIFS access also possible using the same architecture

NFS INDUSTRY CONFERENCE

Scaling on Demand

- Bandwidth and capacity can now be scaled on demand

Need Storage?

Add Disk….

Need Backend Performance?

Scale your Storage Subsystem
Add SAN Connectivity….

Need Higher Network Throughput?

Add a HighRoad Cluster Node….

All while mitigating your risk exposure (it's HA)

NFS INDUSTRY CONFERENCE

Wrap Up and Questions

- NFS Cluster takes us closer to the abstract, virtualized environment where we add what we need, on demand.

- Questions?