



# High Performance Flash Storage Solutions


Burst Buffer Use Cases

**We believe that today's solutions take too long to get the data from where it is stored.**

**That's why we are revolutionizing storage performance to help companies work faster.**



# Traditional Storage Use Cases




## Advantages:

- Isolate storage from server
- Easier to service
- capacity can be re-allocated

## Disadvantages:

- Legacy fabrics (FC/iSCSI)
- constrained BW
- high latency



## Advantages:

- Low latency
- High BW

## Disadvantages:


- Volume size limited to individual SSD capacity
- Fixed capacity per server
- Difficult to service/replace worn out SSDs

# Solution: NVMe over Fabrics

The best of both approaches... plus:

- ❏ Non-Proprietary architecture
- ❏ Standard Ethernet switch fabric
- ❏ Open industry management tools
- ❏ Completely compatible with legacy applications
  - Volumes appear as a local NVMe devices

## Outperforms Today's Fiber Channel & by 10X



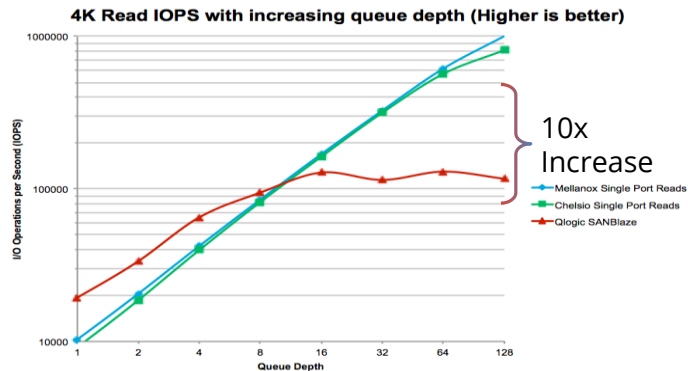
### Advantages of SAN:

- ✓ Isolate storage from server
- ... Easier to service
- ... capacity can be re-allocated

### Advantages of DAS:

- ✓ Low latency
- High BW

## NVMe vs 16Gb Fiber Channel





# NX Series Flash Acceleration Platforms

## Performance/Latency

- ▶ Highest Read/ Write Performance
- ▶ >2.5M RRd/RWr IOPs
- ▶ Consistently low R/W Latency
  - ▶ RD - >200uS at 10 GB/s
  - ▶ WR -> 110uS at 9.8 GB/s



## Technology

- ▶ NX63XX high speed shared memory appliances
- ▶ Shared memory software optimized for low latency networks, utilizing NVMe over Fabrics and pNFS

## Standard NVMe Drivers

- ▶ MSFT, Linux, VMware
- ▶ Industry STD interfaces
- ▶ Management GUI & CLI


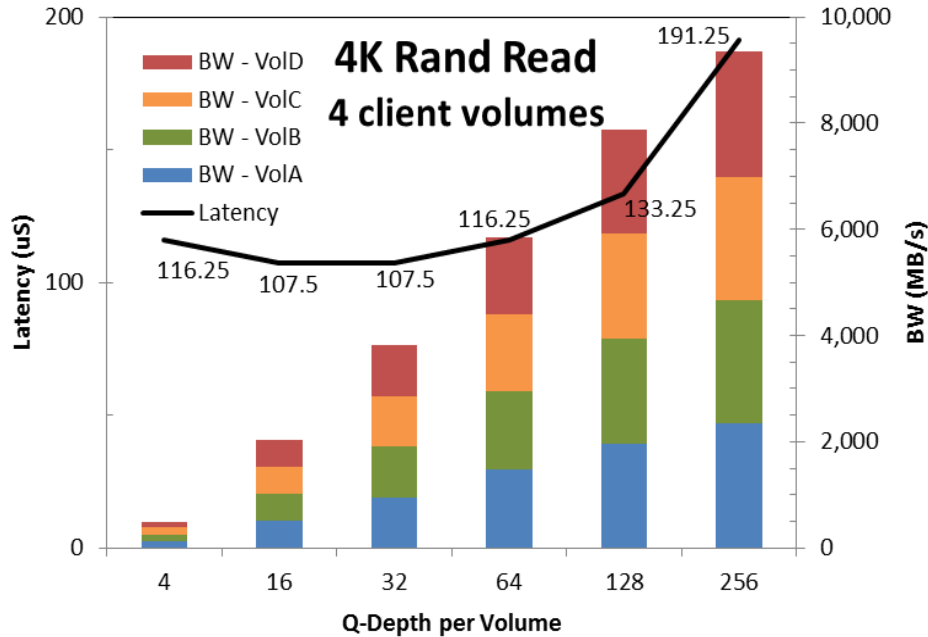


## Capacity/Endurance

- ▶ Std 2U: 8, 16, 32 TB
- ▶ 5 year Warranty, 7 DWPD

## Single System Read Performance ~2.3M IOPs (4K Ran Rd ) < 200uS latency

## Single System Write Performance ~2.2M IOPs (4K Ran Wrt ) < 110uS latency




Notes:  
Identical Q-Depth on each client

# Mangstor HPC Shared Memory Platforms Scales to > 1TB/sec


**Single system:**

**25GBps 4K Random Read at <200us**


**20GBps 4K Random Write at <110us**



Dual EDR  
or  
Dual 100Gbps Ethernet



EDR/FDR/QDR  
or  
10/25/40/56/100 Gbps Ethernet

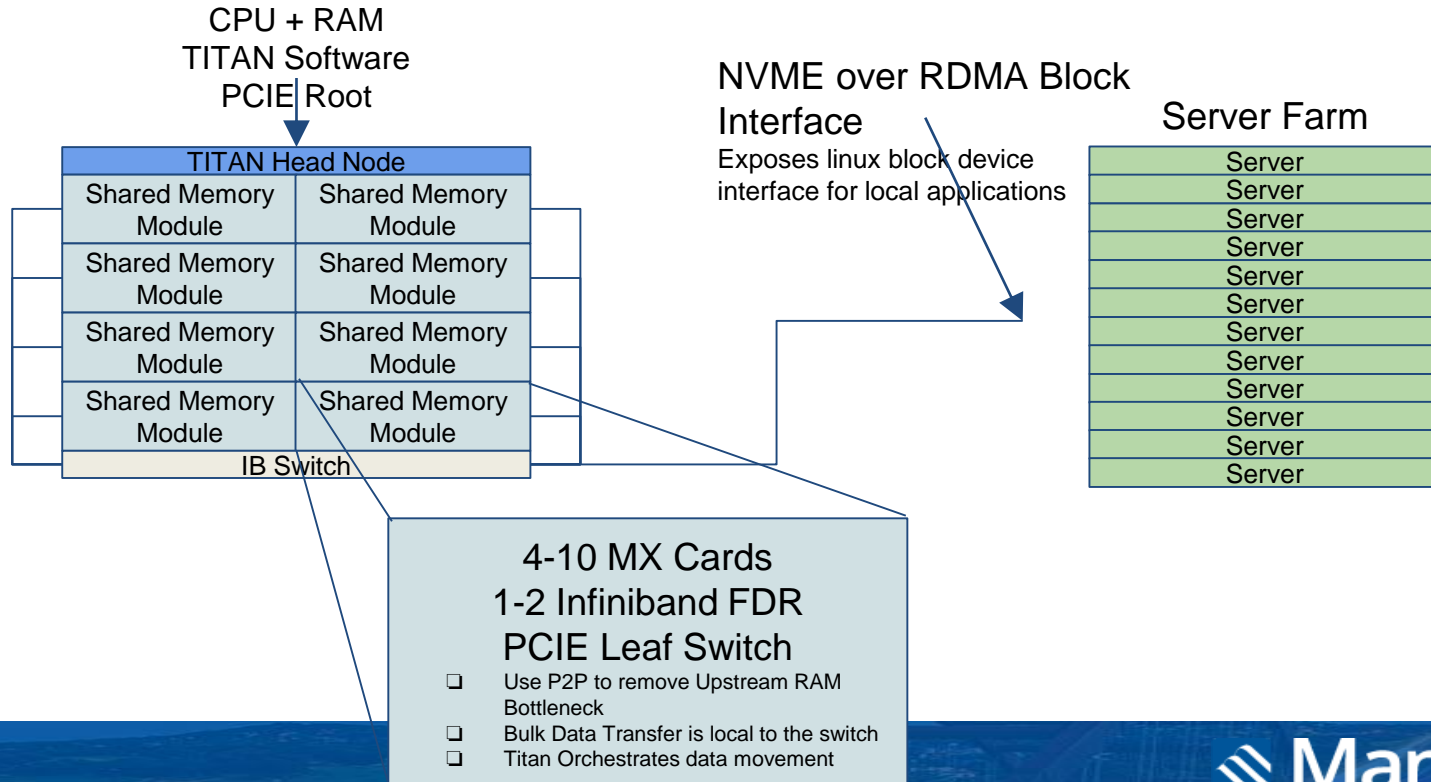


NVME over RDMA Block Interface  
Exposes linux block device interface for local applications

**Volume Manager**

10 MX Cards (25TB-50TB)  
2 Infiniband EDR Connections

# High Performance Bulk Shared Memory - Next Gen





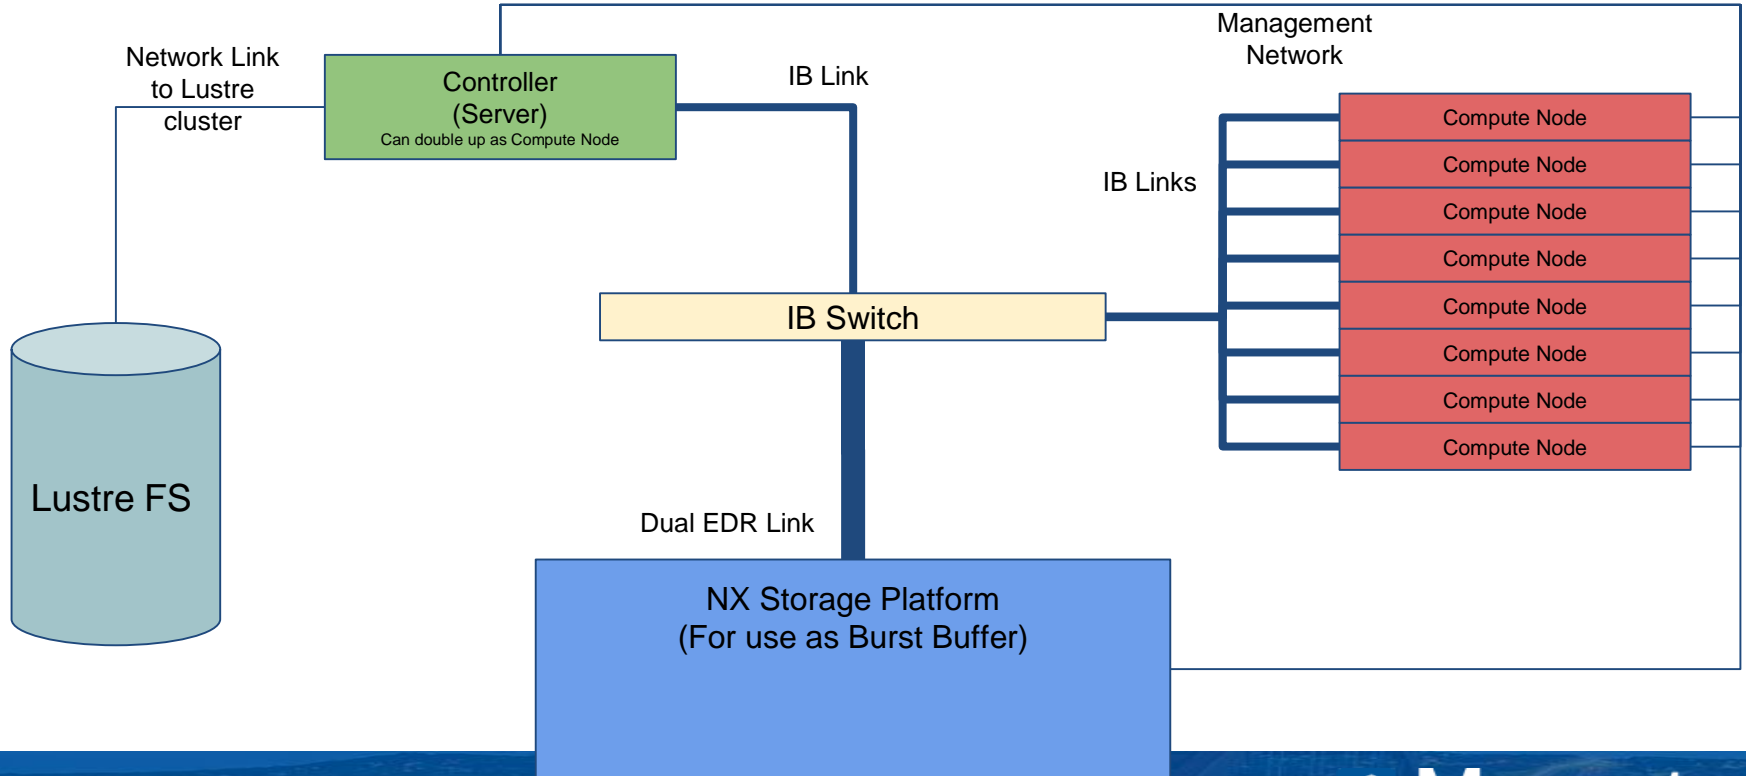
# Mangstor Flash Shared Memory Platforms Use Cases

- ▶ High Speed Burst Buffer
- ▶ High Speed Network Capture
- ▶ 4K Video Capture and Processing
- ▶ High speed and Large Shared Database for EDA Apps
- ▶ High speed write Logging for EDA
- ▶ Intermediate Cache for Data and results for Deep Learning algorithms and other Real Time Data Analysis
- ▶ High speed State snapshot and restore for Large Real Time Data Analysis Farms

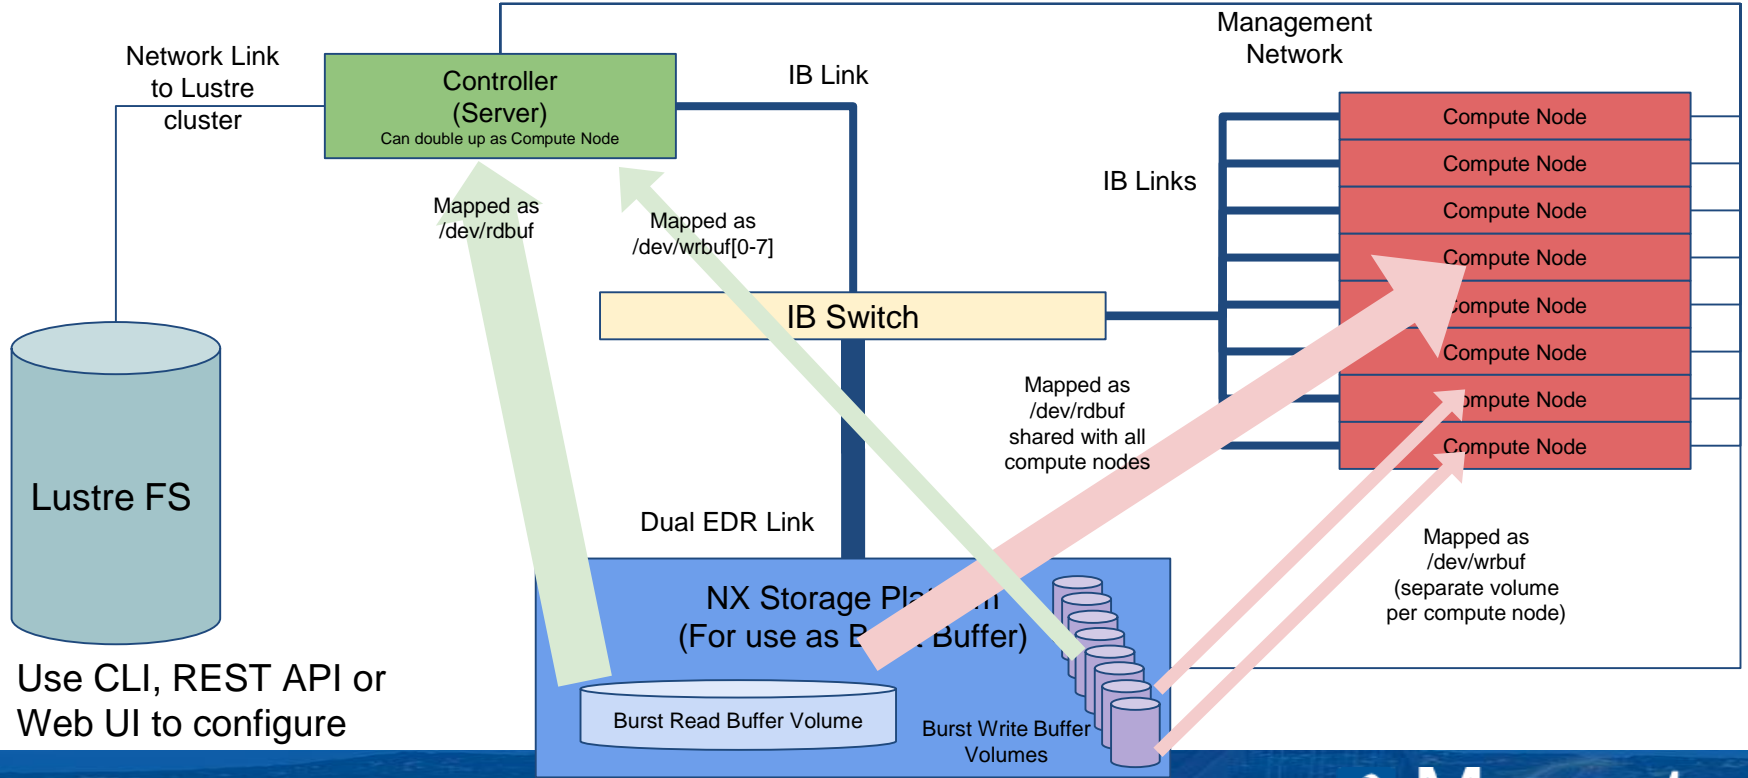
# Mangstor HPC Shared Memory Platform Features

- 25 GBps Aggregate 4K Random Read Bandwidth at < 200us
- 20 GBps Aggregate 4K Random Write Bandwidth at < 110us
- 25TB - 50TB Storage Capacity per unit
- Infiniband, Iwarp and Roce Support
- Mangstor Titan Software for Volume management
- CLI, Web and RESTful interfaces for config and management
- Support for Windows and Linux Servers

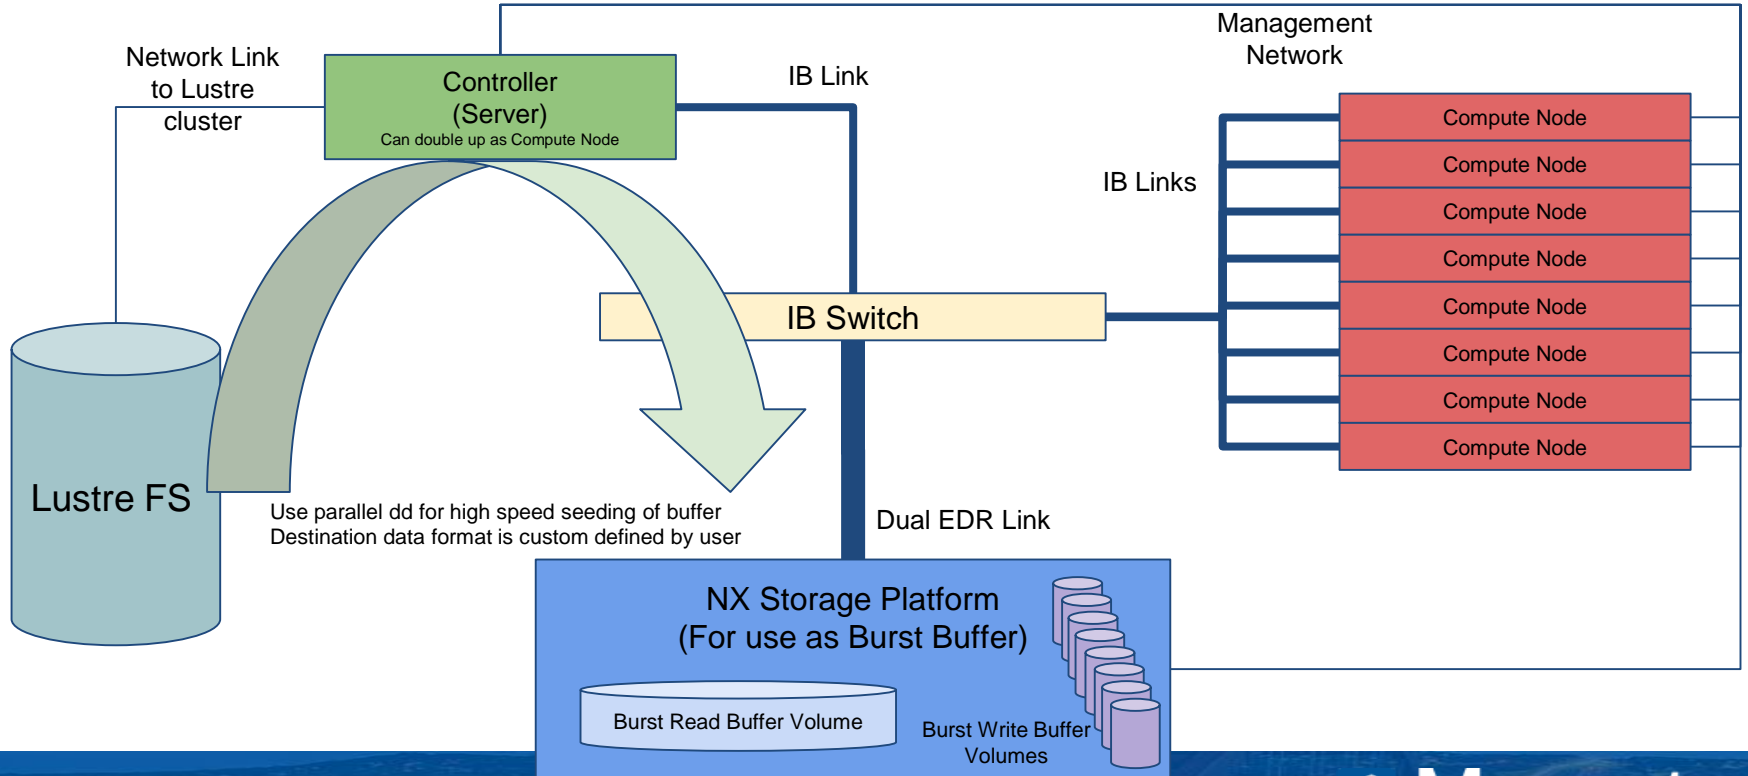
# System Design - Burst buffer with Infiniband links



# Initialization (Time Zero)

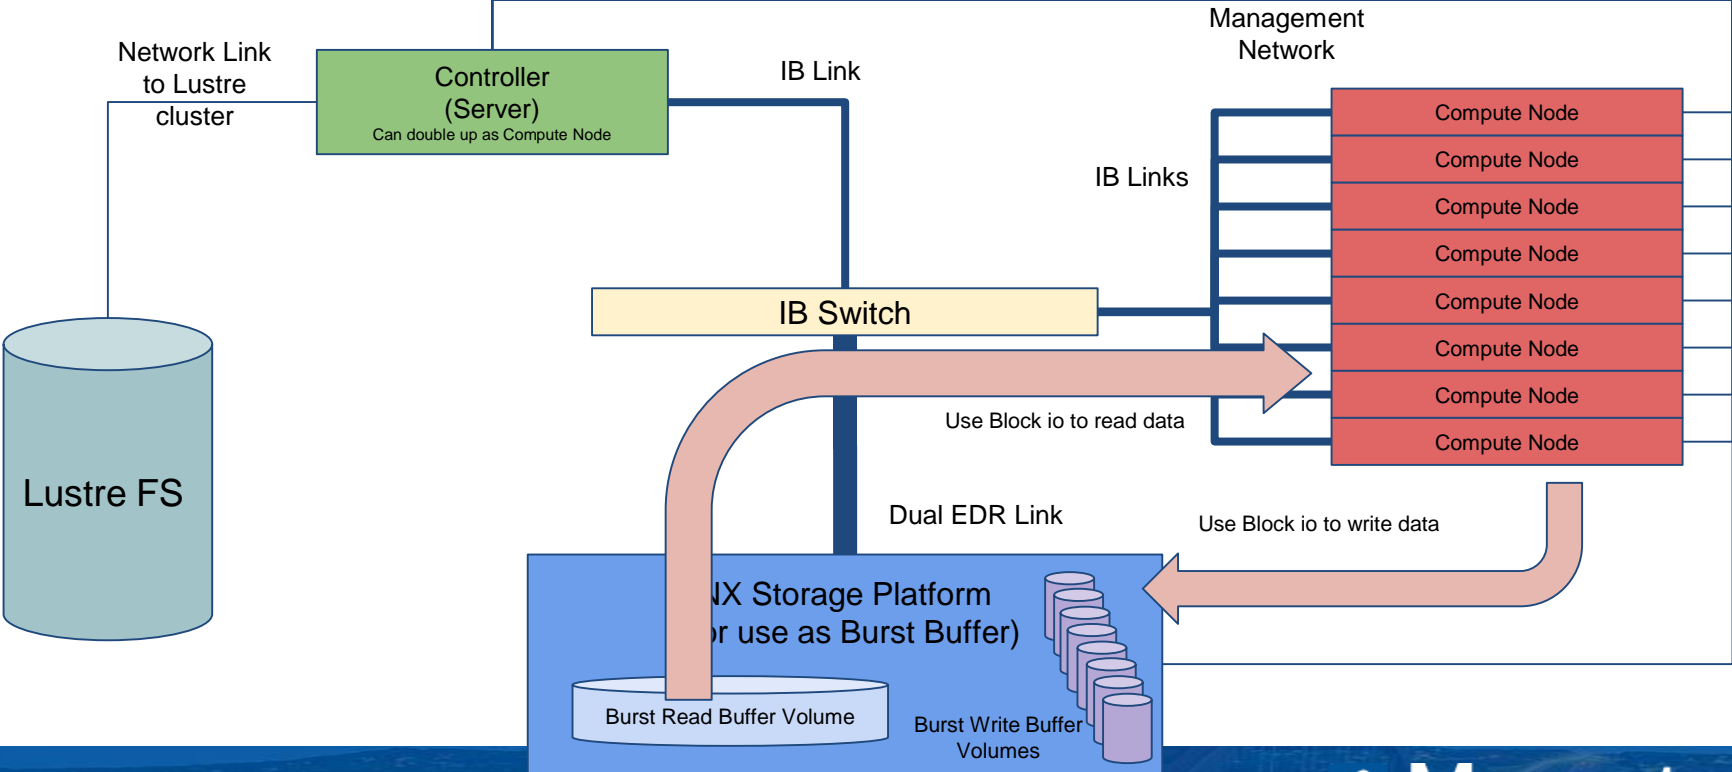


# Seeding Read Buffer

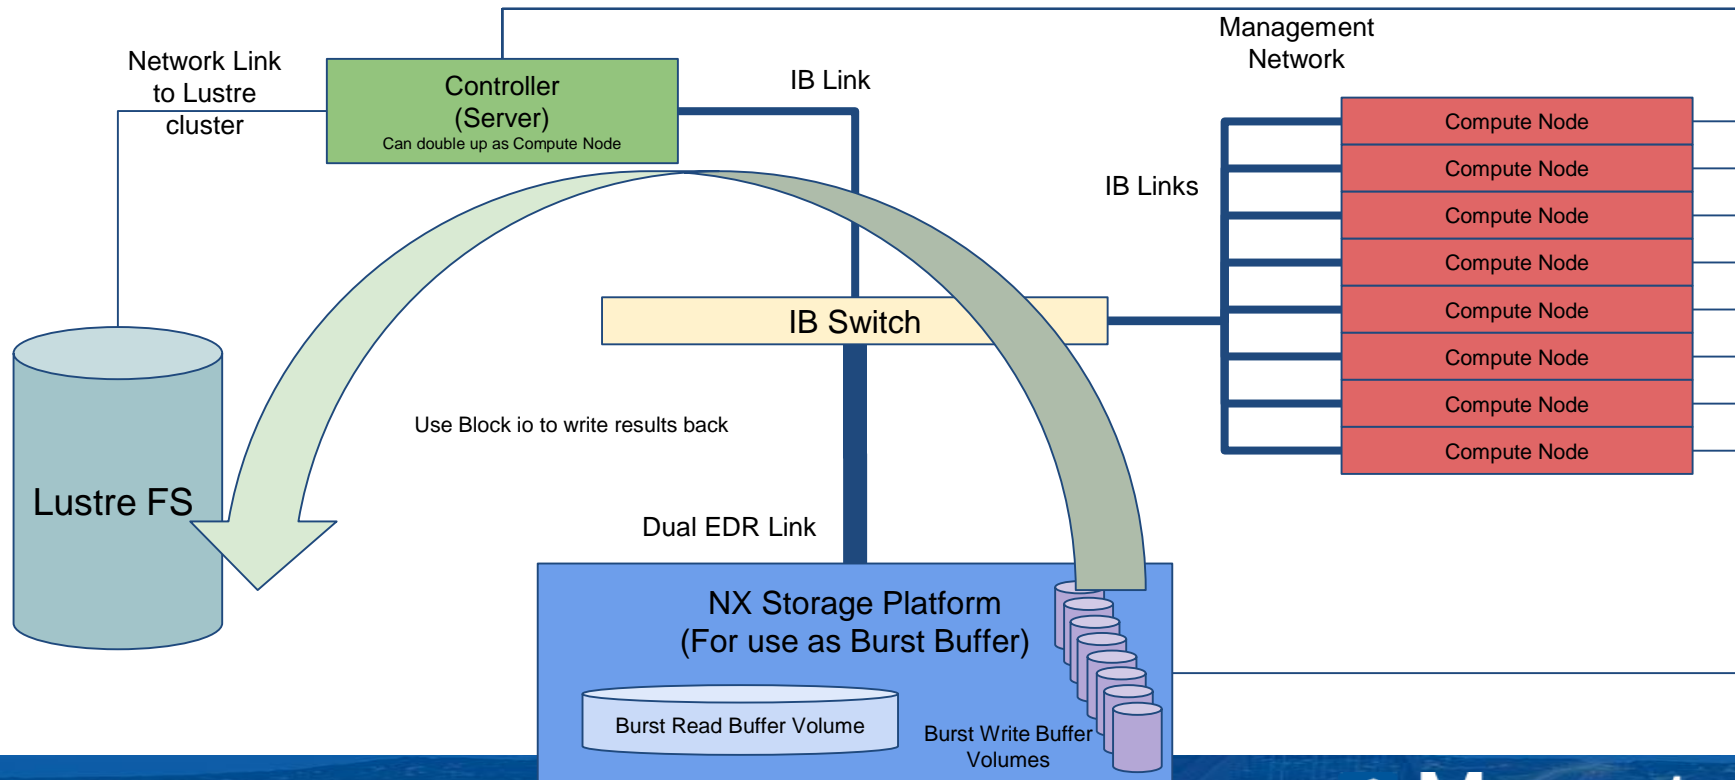




# Compute Node Processing



# Saving Results





**Thank You**