

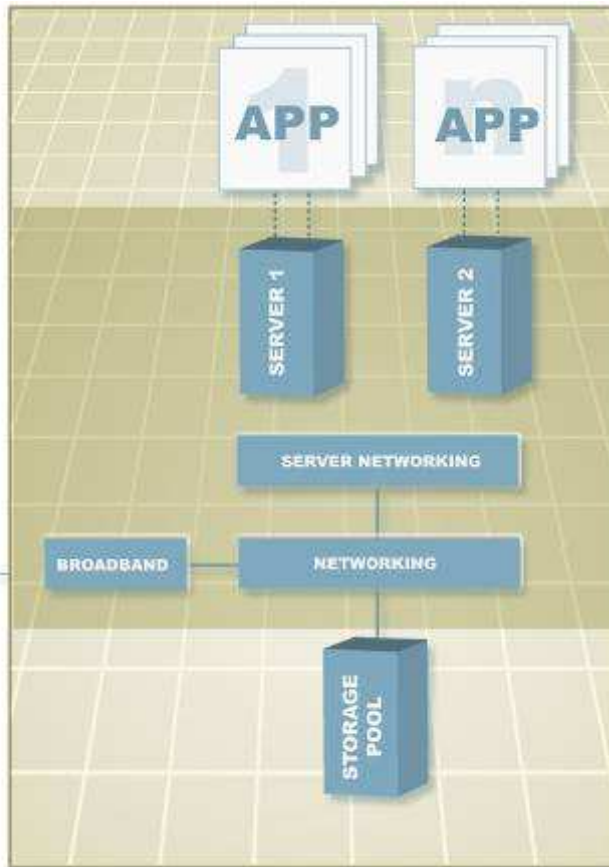
Scalable computing  
without compromise.

LiquidIQ™ and Lustre:  
Converged Communications,  
Parallel File Systems  
and Copious I/O


José Miguel Peleato

# The Corporate Intent

**TODAY**

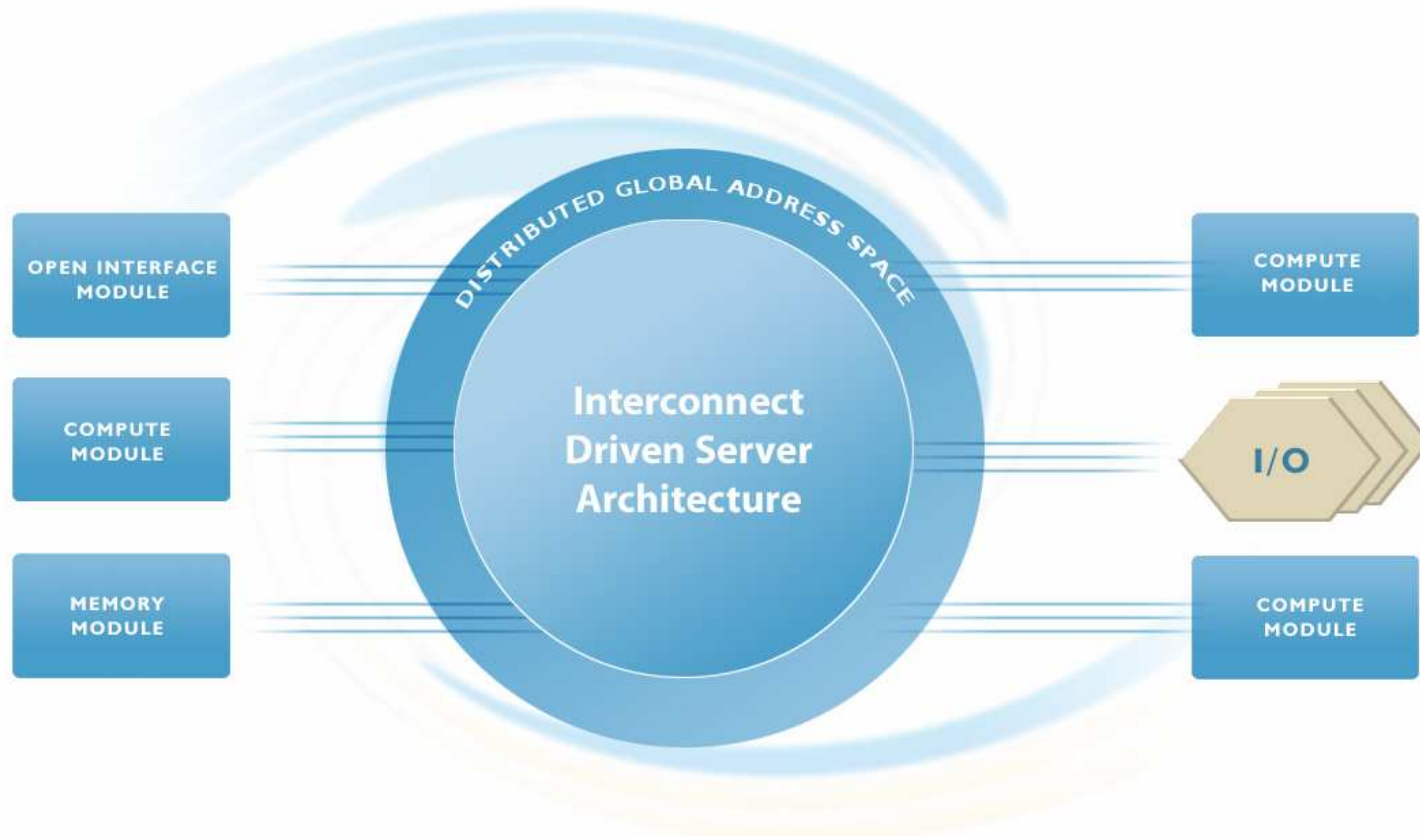


**LIQUID  
COMPUTING**



**The merging of computing and communications to provide a highly dynamic, scalable, and reliable processor pool.**

# Architecture for Scalability & Control



- Performance Optimized Communications and Computing
- Entirely New Control and Configuration Capabilities
- Hot Swappable, Scalable and Self-Managing Components
- COTS computing arranged around a very fast interconnect



# Liquid Computing



ATA VENTURES



VenGrowth

Founded January 2003  
Offices Ottawa, Canada  
Los Altos, California  
Funding Seed Round (May 2004)  
Series A Round (May 2005)  
Status Entering beta program






# LiquidIQ™ System Components

## SERVICE MODULE

- 20 Service Modules per chassis
- Compute Module (CM)
- Memory Module
- Storage Module
- Open Interface Module



## INTERCONNECT MODULE

- 20 Interconnect slots
- I/O Module (IOM)
- Single Chassis Switch Module (SCSM)
- Multi Chassis Switch Module (MCSM)
- Inter Chassis Switch Module (ICSM)



## CHASSIS

- Power Entry Module (PEM)
- DC or AC Powered
- Mid-Plane (MP)




## CHASSIS MODULE


- Alarm Interface Module (AIM)
- Alarm Module (AM)
- Network Synchronization Module (NSM)
- Synchronization Interface Module (SIM)
- Fan-Tray Module (FTM)




IOM



SCSM



ICSM



MCSM





## LiquidIQ™ : A Superior Platform for Lustre - I

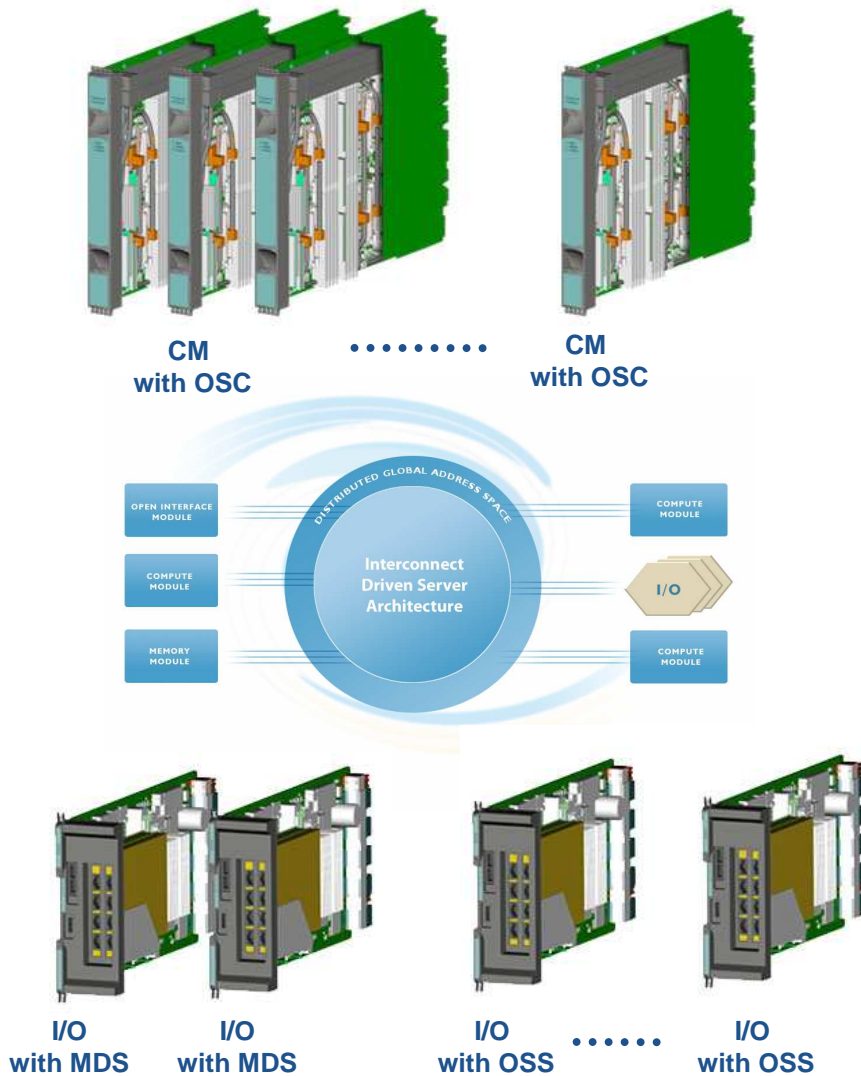
- Communications driven architecture
  - Direct Hypertransport connection from each processor to I/O. No PCI bus funnel
  - Modular multi-plane communications
  - Cross Bisectonal Bandwidth = 1
  - Point to point links
  - Computing and I/O node location independence
- Copious I/O
  - 2.4 Tb/s I/O per system (12 chassis)
  - PCI Express 1 GE, 10 GE, 2 and 4 Gb/s Fiber Channel interfaces



## LiquidIQ™ : A Superior Platform for Lustre - II

- Copious Interconnect
  - Low latency and high bandwidth (16 GB/s)
- Multi-plane Inter-processor Communications (IPC)
  - Traffic bursts are spread over communications planes
  - Load balancing
- Reliability
  - Multi-plane communications provides IPC redundancy
  - IO port back-up
- Future proof
- Lustre configuration flexibility


# High Performance LiquidIQ™ Lustre Integration



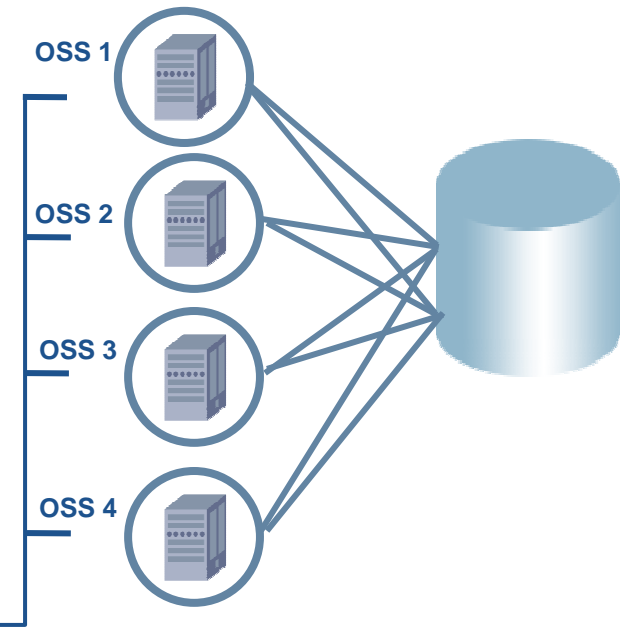
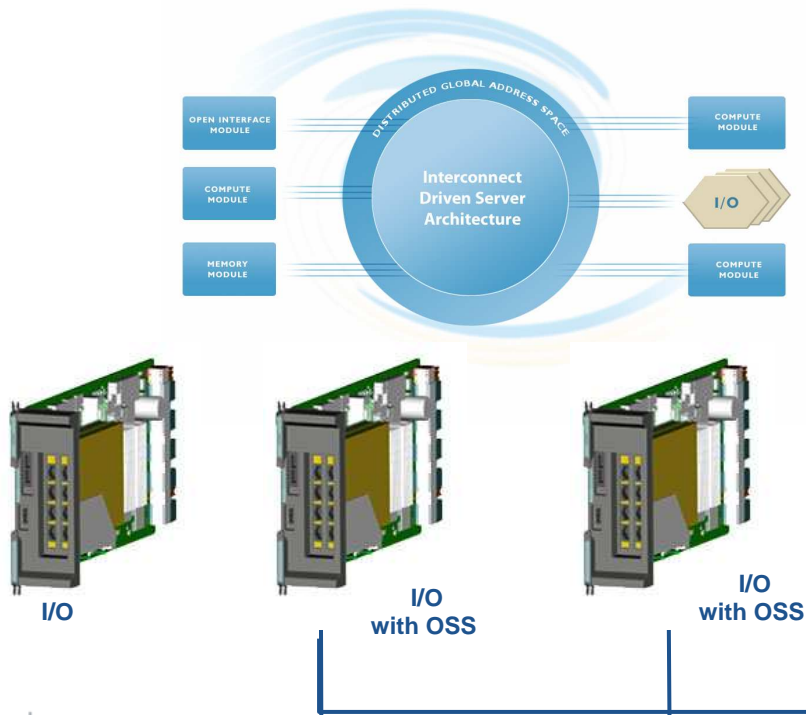
- OSC in Compute Nodes
- OST and MDS in I/O Modules
- Self-contained
- Fast internal communications > High performance



# LiquidIQ™ Lustre High Performance MDS



- MDSs can be installed in CMs (8 way Opteron).
  - High performance
- OST can be external or in I/O Modules

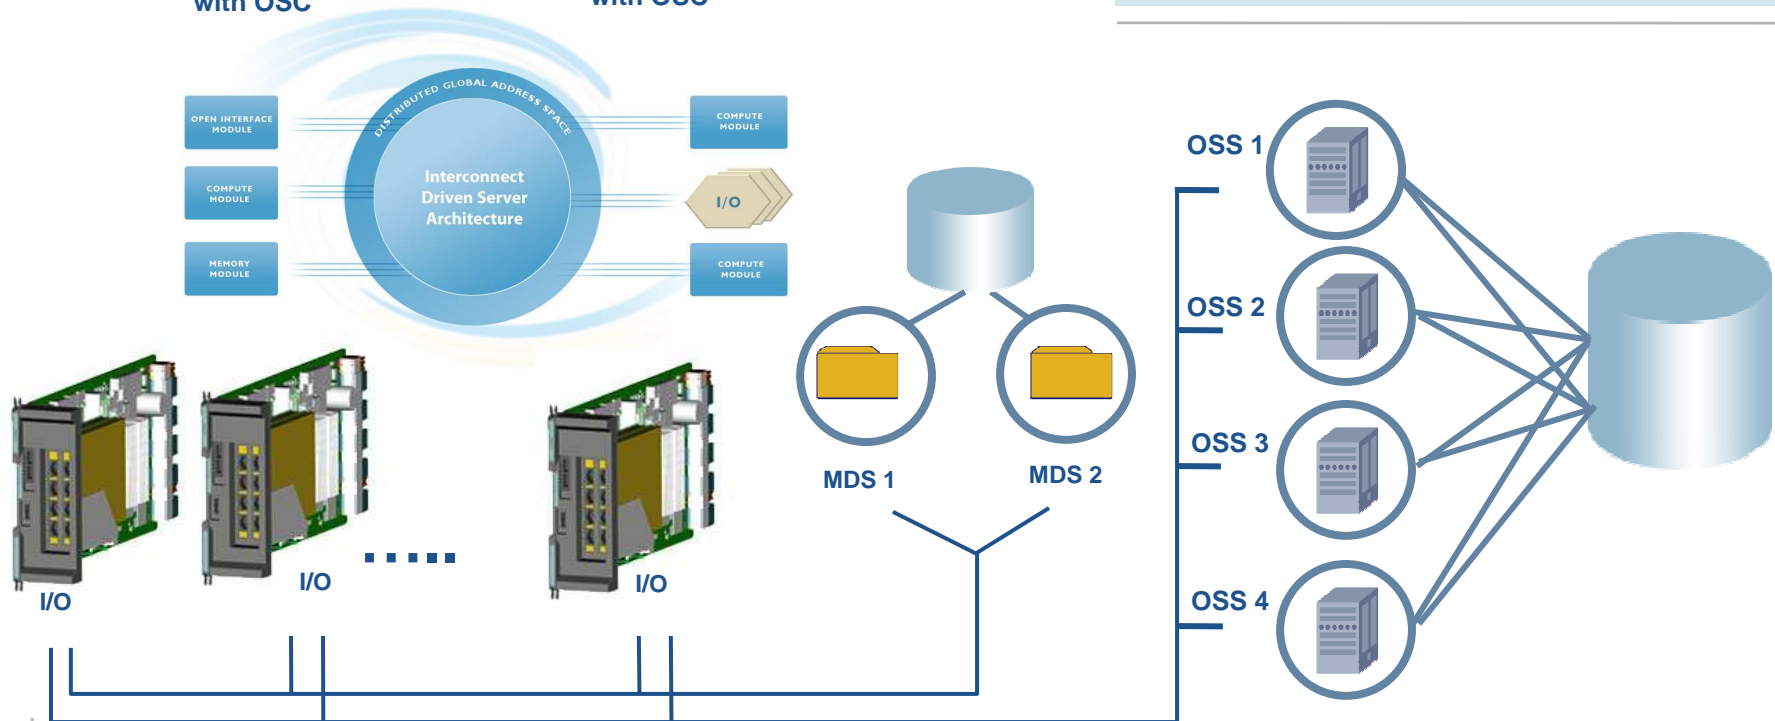




# LiquidIQ™ Lustre integration into existing Lustre infrastructure



- LiquidIQ integration into existing Lustre infrastructure
- OSC in CMs
- IOM connects to MDS and OSS





## Availability

- LiquidIQ entering beta phase in May
- Lustre LiquidIQ beta in July with Lustre 1.6
- LiquidIQ GA in Q3 2006