



**8.12.2016**

# **Kinetic at CERN**

openlab Workshop

Paul Hermann Lensing



SEAGATE



**CERN**openlab

- › **Hard drives... but**
  - **Ethernet instead of SATA**
  - **Key-Value instead of block interface**



1 PB CERN Kinetic Installation



- › **TCO** no need to buy, power, or administrate servers to provide network accessible storage
- › **Robustness** replicating key-value functionality on each drive isolates failures and performance bottlenecks
- › **Flexibility** add capacity / replace a failed drive anywhere in the network
- › **Abstract Interface** key-value API hides host-aware technology such as Shingled Magnetic Recording from the user / application



# KINETIC

## Open Storage Project

Platinum Members

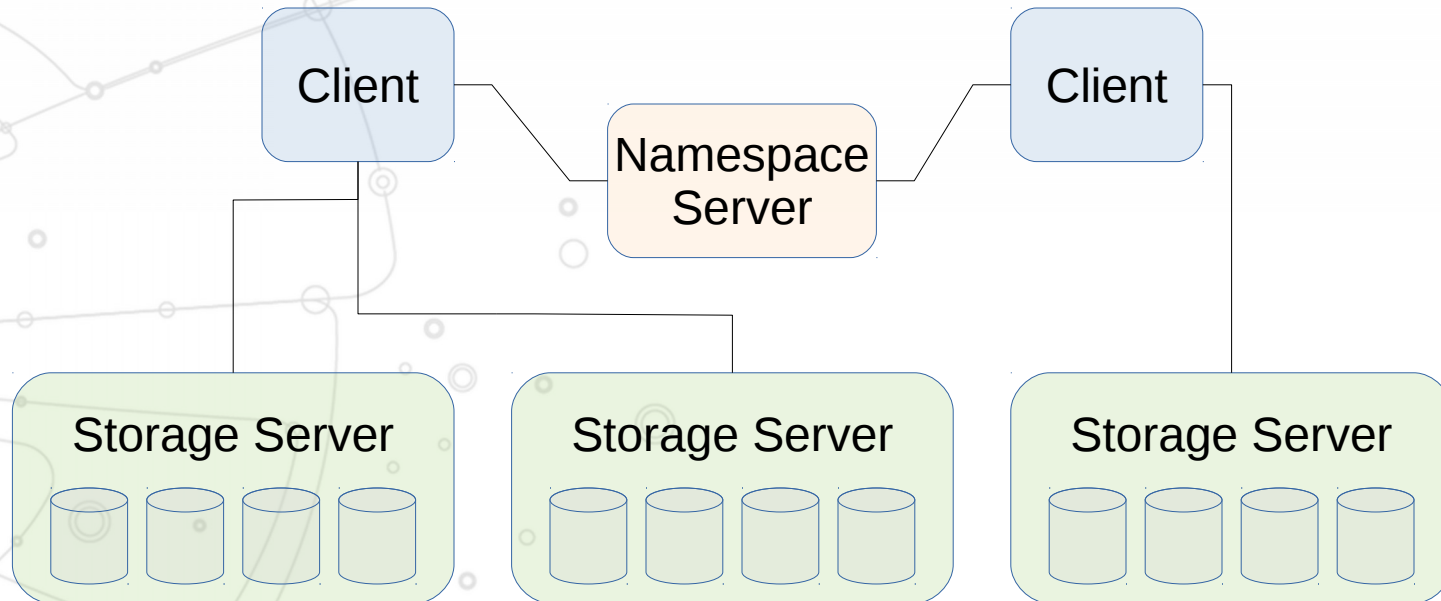


<http://www.openkinetic.org>

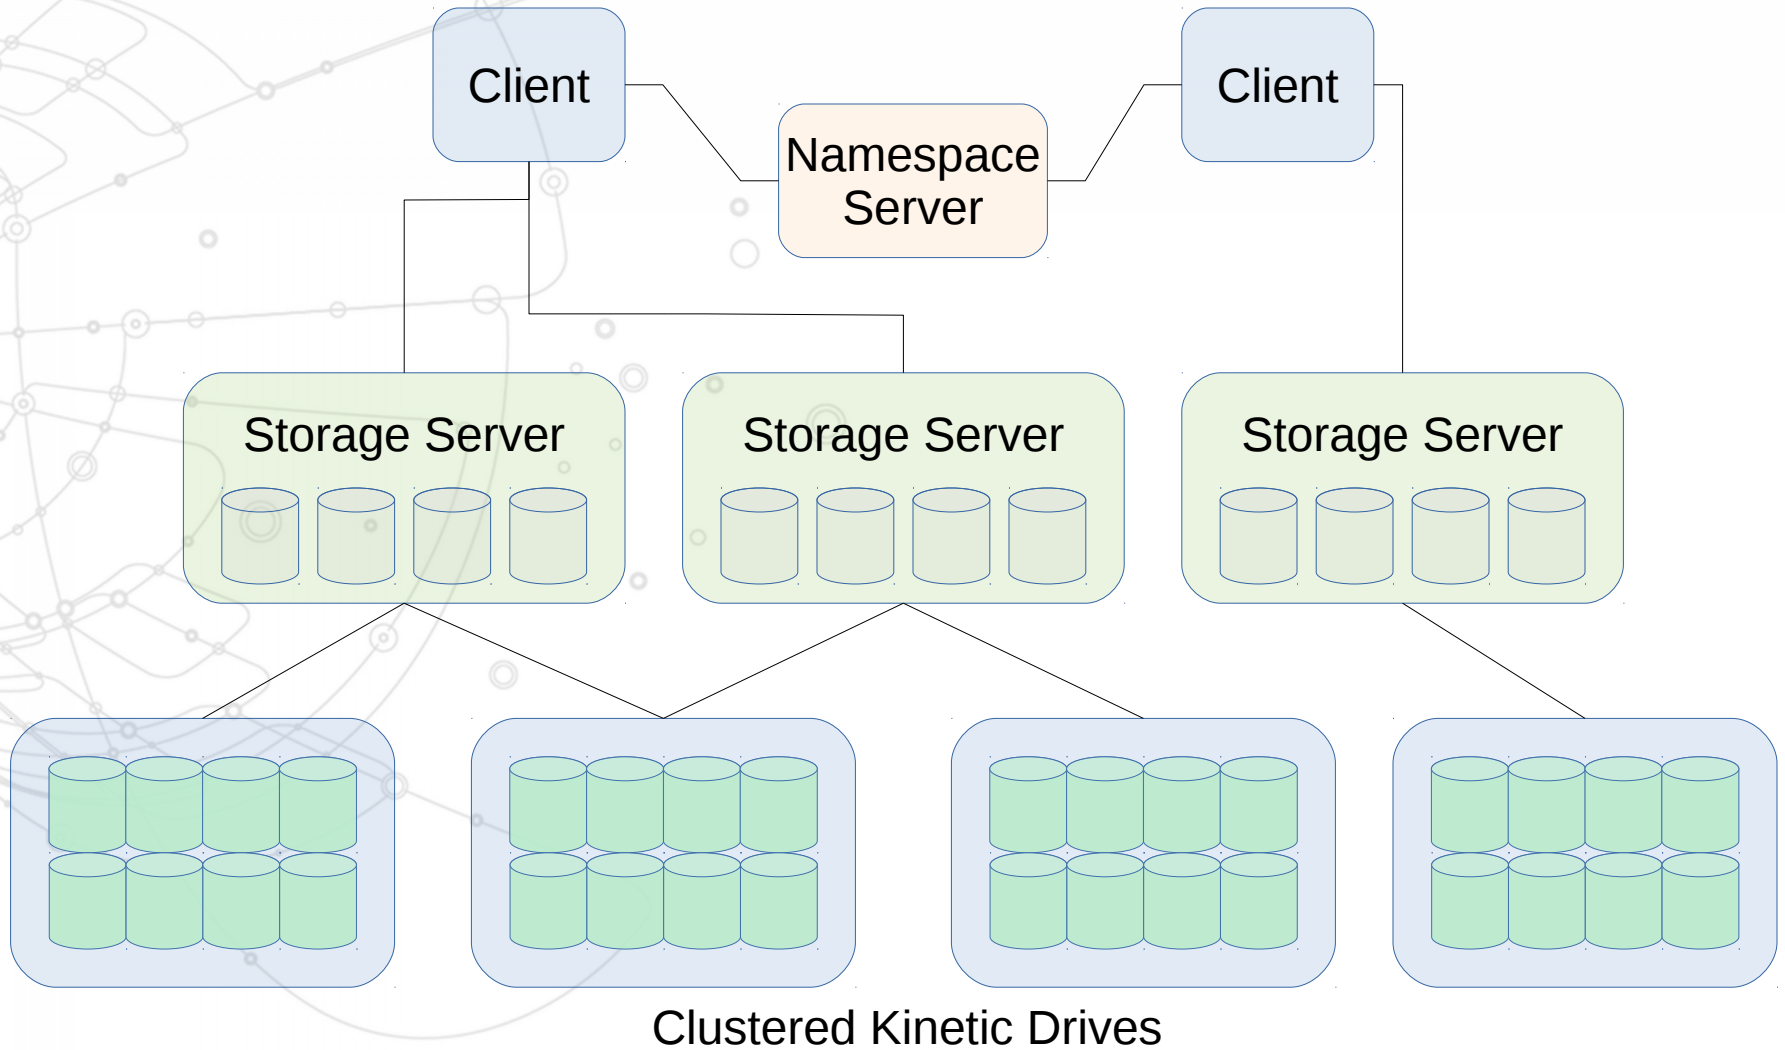
Silver Members



# Basic EOS Architecture



# EOS Extended with Kinetic



# Clustering Kinetic Drives

- › **Simplify EOS Development:**
  - **Can treat a kinetic cluster as a single, highly reliable, very large disk**
- › **Simplify EOS Operation:**

```
#-----#
# type #      name #  groupsize #  groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #
#-----#
spaceview  default  0          0          97         97         263.69 T      385.89 T      385.89 T
spaceview  kinetic  0          0          4          4          322.12 T      400.00 T      400.00 T
```

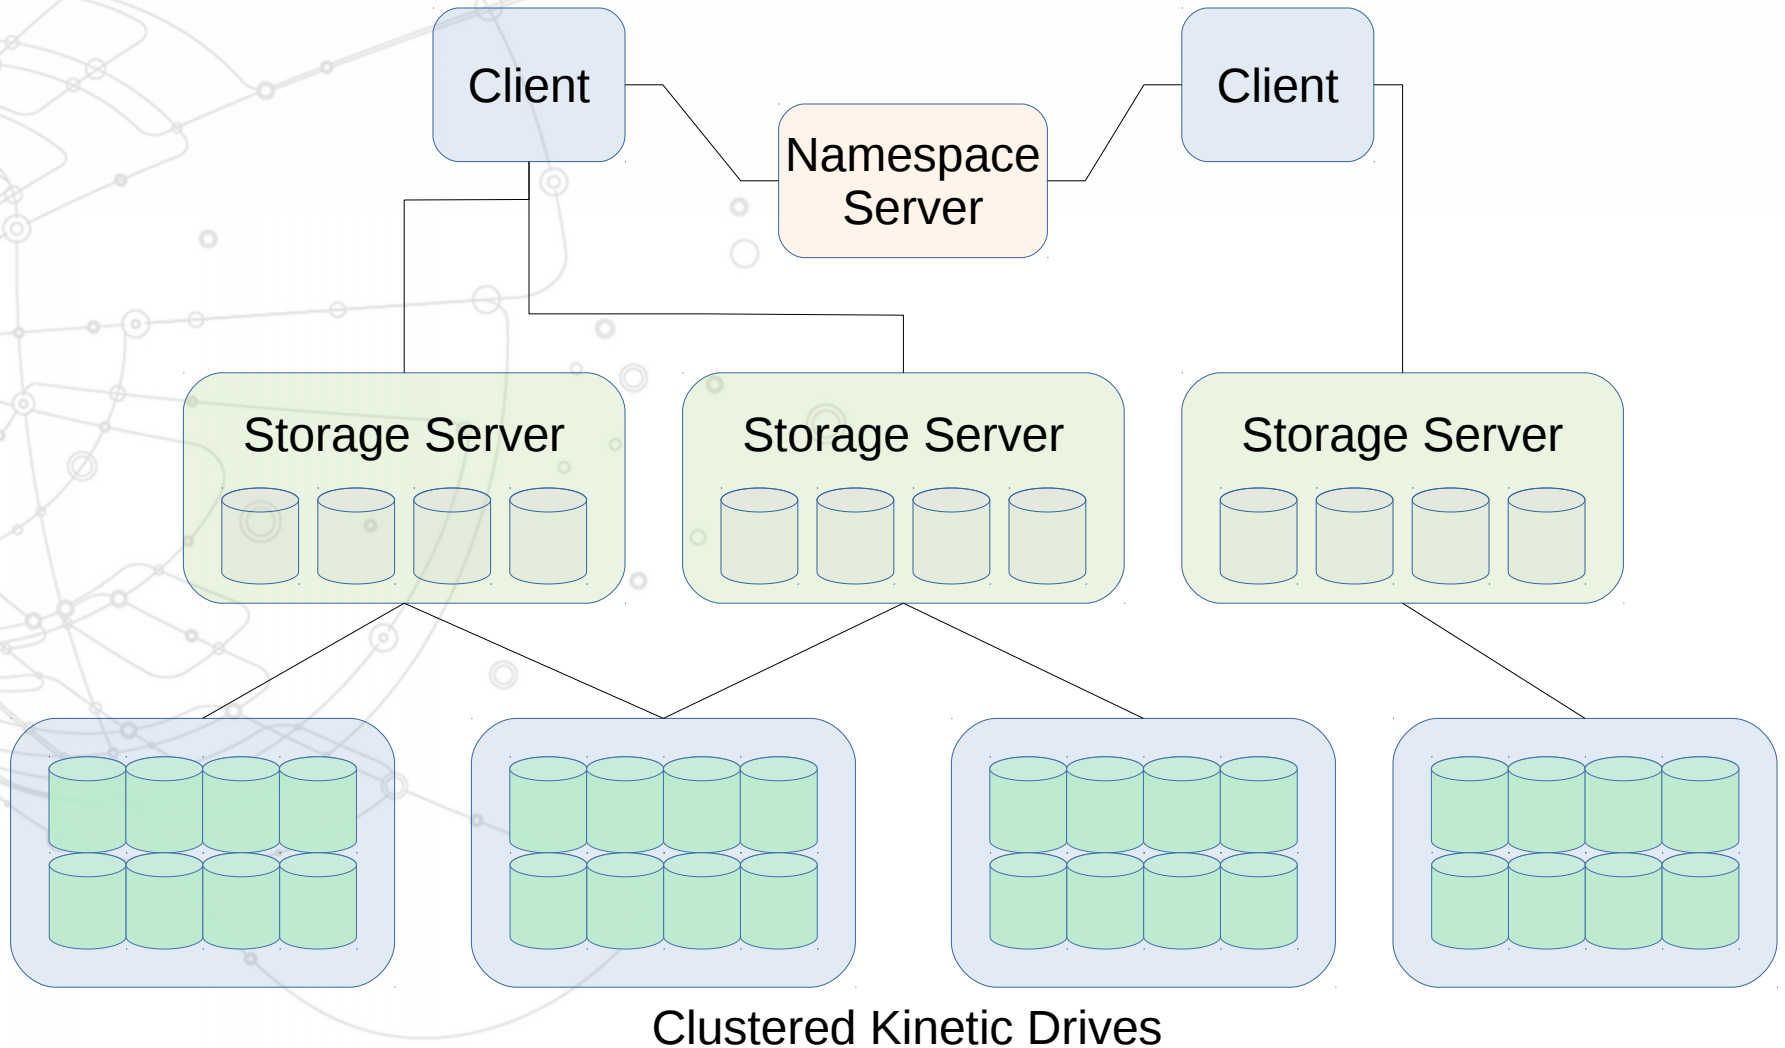
- › **Optimized redundancy handling: erasure coding and replication on key-value data chunk level**
- › **Improved single-stream performance**
- › **Health monitoring / repair on cluster basis**

# EOS with multi-path and Kinetic

- **Current EOS deployment:**
  - > **1500 EOS storage servers**
  - **even at peak < 5% utilization of combined storage server network capacity**
- **Multiple data paths allows any subset of clusters to remain available at high bandwidth regardless of current utilization**
  - **only one data path to locally attached storage**
- **An existing group of EOS storage servers can supply (much) higher storage capacity => save \$\$\$**



# EOS Extended with Kinetic



- **2<sup>nd</sup> generation Kinetic drives @ CERN**
  - **2 PB of additional Kinetic Storage**
  - **mature technology**
  - **higher capacity / performance**
- **Kinetic integration is stable and merged into EOS master branch**
- **Main goal: get Kinetic storage into full production use**