

# Managing a High Availability Lustre Environment Using Multiple Namespaces

April 30th, 2024  
Brad Gipson

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

# What is NCCS?

- National Center of Computational Sciences operated out of Oak Ridge National Laboratory (ORNL)
- ORNL has 9 directorates with many science research organizations, one of which is CCS (Computing and Computational Sciences) which NCCS is operated under
- NCCS supports many programs from many organizations including:
  - Department of Energy (DOE)
    - ORNL Leadership Computing Facility (OLCF)
  - National Oceanic and Atmospheric Administration (NOAA)
    - National Center for Computational Sciences (NCRC)
  - Air Force Weather
  - Many others
- Not a service provider, all partnerships are research-based



# Lustre Hardware Overview

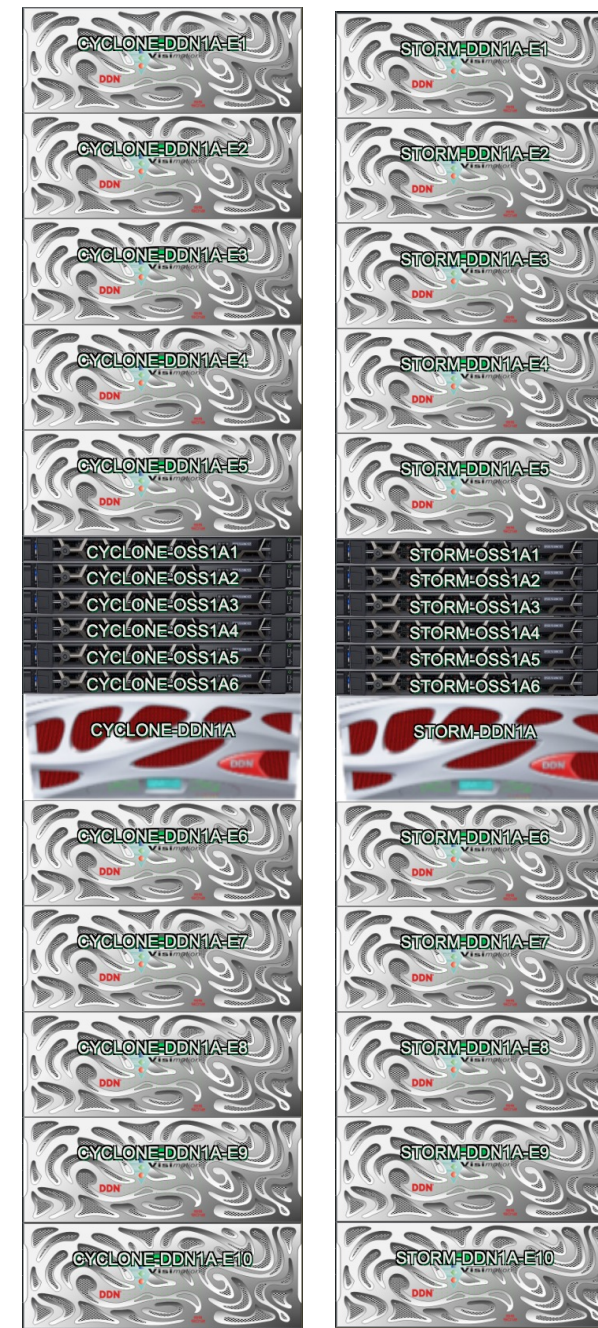

- Two Separate Filesystems
    - Storm
    - Cyclone
  - Lustre 2.15 (b2\_15)
  - Matching TDS
    - Gale
- Each Filesystem Contains:
    - Capacity:
      - 7 PB Per FS
      - DDN SFA14KX
      - 840x 12 TB PMR HDD
    - Metadata:
      - 110 TB Per FS
      - DDN SFA200NV
      - 20x 7 TB 1DWPD NVMe
    - Network
      - 2x100GB/s HSN





# Physical Attributes

- Per Filesystem
  - 2 MDS nodes (1x MDT per node)
  - 6 OSS nodes (2x HDD OSTs per node)
  - 1 MGS node (1x per namespace)
- 2 management servers
- 2 Utility Servers
- Misc. hot spare nodes



# HA Overview

- Two separate namespaces
  - Storm/Cyclone
- Identical in size/shape/configuration
- Redundant management stack w/ failover
- Replication between namespaces
- Dual homed Lustre hosts



# Design Considerations

Ensure that both namespaces are same capacity

Same/Similar configurations eases upgrade paths

Environment aware user-level replication

Local data transfer nodes

Redundant high speed network links  
(Optional/Recommended)

# Workflow Management

Users write their jobs to  
/lustre/active

This directory links groups to  
directories on either Storm or  
Cyclone based on GroupId  
or filesystem status.

Groups are gathered and  
“assigned” a filesystem  
under normal operations

Normal Operations: /lustre/active

- GroupA -> /lustre/storm/GroupA
- GroupB -> /lustre/cyclone/GroupB

Storm Maintenance: /lustre/active

- GroupA -> /lustre/cyclone/GroupA
- GroupB -> /lustre/cyclone/GroupB

Cyclone Maintenance: /lustre/active

- GroupA -> /lustre/storm/GroupA
- GroupB -> /lustre/storm/GroupB

# Workflow Management

- Initial job data is transferred into both namespaces from external sources
- Compute jobs choose which data to run on and where to export to based on /lustre/active
- Job output data is then replicated between filesystems and exported externally
- In the case of filesystem maintenance, replications are queued until second namespace is restored



# Pros/Cons

## Pros

- Namespace availability
- Project-level replication
- Non-blocking updates
- Early issue identification on opposite namespace

## Cons

- Cost
- Imbalance of user I/O bandwidth
- Technical complexity
- Post-outage data replication “buildup”

# Recent Additions



Lustre 2.15 upgrade



Data on Metadata



Automatic Failover of  
OSS/MDS nodes (Soon!)

# Summary

TWO SEPARATE BUT  
IDENTICAL  
NAMESPACES

USERS ACCESS DATA  
USING LINKS IN  
/LUSTRE/ACTIVE

CHANGE  
/LUSTRE/ACTIVE LINKS  
TO WORK ON  
NAMESPACE WITHOUT  
INTERRUPTION

USERS SEND DATA  
INTO BOTH  
NAMESPACES AND  
REPLICATE ON JOB  
COMPLETION

# Discussion/Questions

[gipsonbm@ornl.gov](mailto:gipsonbm@ornl.gov)

