

# Adding HSM to existing Lustre FS

RUG 2016 – Paris

([dtacchella@cray.com](mailto:dtacchella@cray.com))

## Topics

- **Initial status**
- **Challenges**
- **Encountered issues**

The Cray logo is positioned in the top right corner, featuring the word "CRAY" in a blue, sans-serif font. To the right of the logo is a decorative graphic consisting of a grid of small circles, some of which are colored in shades of blue, red, and green, creating a pattern that resembles a stylized molecular structure or a data visualization.

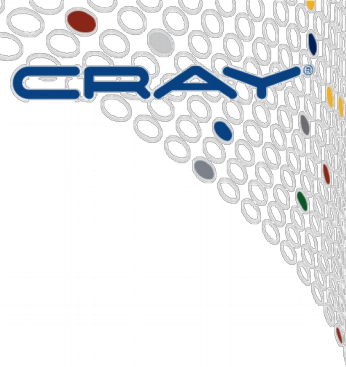
CRAY

---

COMPUTE

| STORE

| ANALYZE

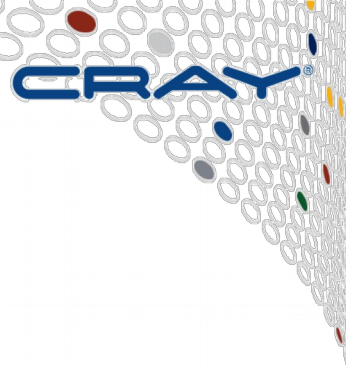


Initial status - boundaries

## Filesystem with

- **460 million inodes**
- **7.71 PB of data**
- **Robinhood run in scan mode**

## Challenges 1/2



- **Robinhood in scan mode**

- **Takes few days to scan whole filesystem**

- **Directory move + file remove caused GC to fail**

- **Modify “one\_path” and “this\_path” functions to consider only one result**

- **Migrate files by volume and file count**

- **i.e : 10M files and 100TB of data**



- **Lustre Action queue**

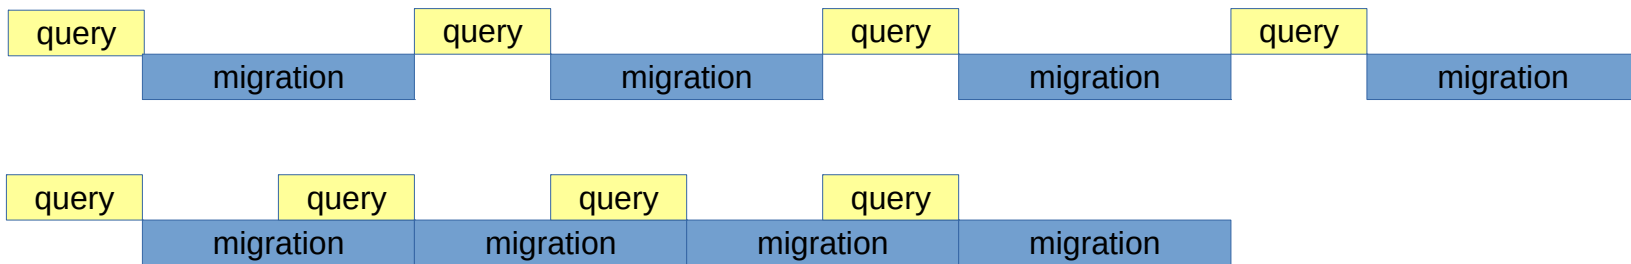
- **Responsiveness of AQ rapidly decreases into a dead-lock situation as queue size increases (LU-7988)**

- **Use external tool for file archival, keeping track of current AQ size**



## External migration tool

- Custom developed
- Injection bandwidth and entries can be tuned any time
- File pickup by size and/or last modify
- Query/Migration overlap (discard duplicate entries)



COMPUTE

STORE

ANALYZE

# Questions ?