

Staging ~2 Million Files from Tape for a User

Wednesday, September 11, 2024 11:30 AM (20 minutes)

Imagine you turn on your work laptop or arrive at the office and find this message from the customer service team: “we have a user that is trying to retrieve *many* files from the HPSS archive. At the current retrieval rate, we estimate it will take 6 months for the user to retrieve all the files in the dataset. Can you help?”

What do you do? How do you proceed? What features does HPSS offer to help with this situation?

I’ll answer those questions and more as we examine LLNL’s approach to retrieving nearly 2 million files across 10’s of tape volumes with tools like quaid, SQLite, and RabbitMQ (along with a bit of custom Python code).

Timeslot

15 min

Author: Mr CLEARY, Geoff (LLNL)

Presenter: Mr CLEARY, Geoff (LLNL)