CTA Backpressure

Current status

PRESENTED BY CEDRIC CAFFY

Plan

1. Backpressure
    1. Definition
2. Retrieve request with backpressure
    1. Setting up disk system
    2. Queueing of the Retrieve Request
    3. Scheduler getNextMount()
    4. Mount getNextJobBatch()
3. Summary
4. Questions

1. Backpressure

1. Definition
   - Mechanism that ensures that a certain amount of free space is available in an EOS space

An EOS space can have multiple filesystems !

2. Retrieve request with backpressure

1. Setting up a disk system
   - cta-admin command
   - Disk system
     - Allows to query the free space of an EOS space
   - What defines a disk system
     - A unique name
     - A free space query URL: example: *eos:eos_instance:name_of_eos_space*
     - A regex to match a disk system with the destination URL of a file
     - A refresh interval: how long the queried free space will be used?
     - A targeted free space: how much free space do we target in the EOS space ?
     - A sleep time: how much time the queue should sleep when the targeted free space is reach ?

2. Retrieve request with backpressure

2. Queueing of a Retrieve request
   - processPREPARE()
     - Get the destination URL from EOS

       Example: root://eos_instance//eos/ctaeos/preprod/directory/file_name?eos.lfn=fxid:83f&...&eos.space=default
   - scheduler.queueRetrieve()
     - Get the disk system name according to the file destination URL (regex matching)
     - Create the Retrieve Request and assign the matched disk system to it
     - Queue the Retrieve Request

2. Retrieve request with backpressure

3. scheduler getNextMount()
   - No mount is returned if the queue is sleeping

2. Retrieve request with backpressure

4. Mount getNextJobBatch()

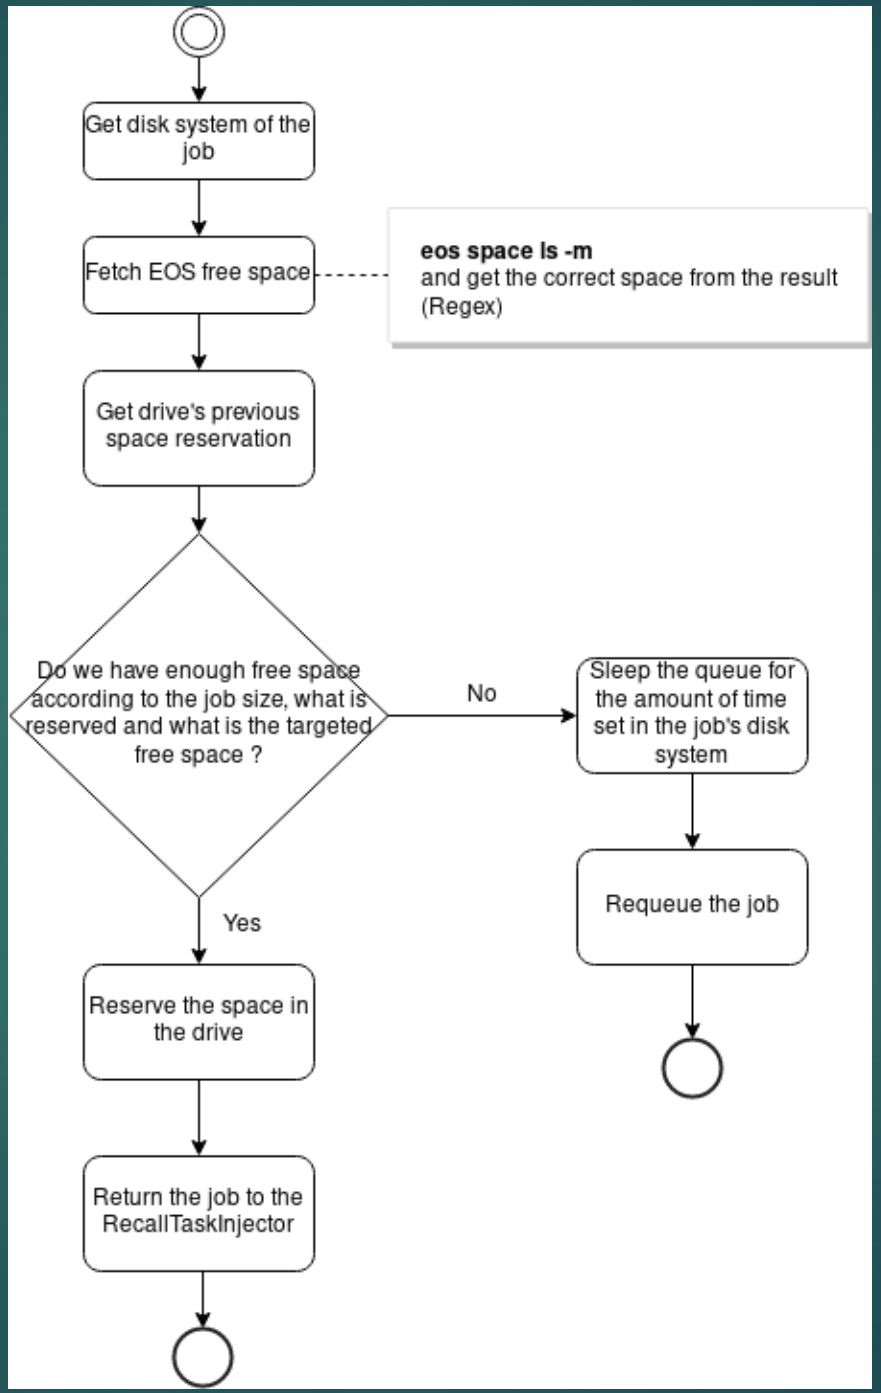
Get disk system of the job
Fetch EOS free space
eos space ls -m
and get the correct space from the result (Regex)
Get drive's previous space reservation
Do we have enough free space according to the job size, what is reserved and what is the targeted free space ?
No
Sleep the queue for the amount of time set in the job's disk system
Requeue the job
Yes
Reserve the space in the drive
Return the job to the RecallTaskInjector

3. Summary

- Backpressure only for Retrieves
- Disk system concept
  - Fetch EOS free space
  - Drive reserve space
  - Queue sleep time
- If no files can be written to the EOS disk (job size is too big to comply with the targeted free space)
  - Sleep the queue for X amount of time

4. Questions

- As a space can contain multiple filesystem
  - Two disks, one full and the other empty
    - eos space ls –m will say that we have enough space, but what happens if the retrieved file goes to the full disk ?