

Intra-queue Fair-share Preemption Policy

Requirements:

1. In intra-queue preemption, we want to support preemption based on FairOrderPolicy along with FifoOrderPolicy. This is to ensure that no job gets starved and all jobs are able to get equal resources.
2. All users in a queue should be getting equal resources.
3. All apps for a user should be getting equal resources.
4. Users should maintain their user-limits after preemption.

Terminology:

Term	Formula	Description
Fair share per app	$\frac{\text{User Limit of app's user}}{\text{number of apps of that user}}$	UserLimit of the user of the app divided equally among all apps belonging to that user
User-Limit	$\max \left\{ \frac{\text{total queue resources}}{\text{number of users}}, \text{user-limit-percent} \right\}$	The max percent of resources a user can use.

Design:

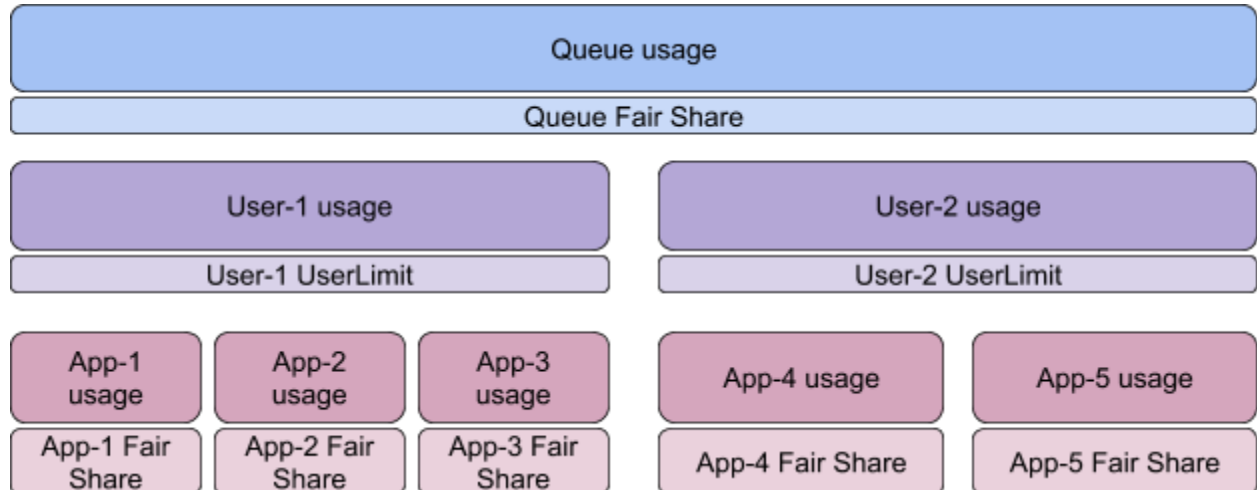
For fair share preemption, we check for following 2 conditions:

1. After preemption, users consume utmost upto their userLimit, if there are pending requests in the queue.
2. After preemption, an app (whose container is being preempted) doesn't drop below its fair-share.

Fair-share is calculated for each app of a user.

For an app, fair share is calculated as: $\frac{\text{UserLimit of the user of that app}}{\text{number of apps of that user}}$, i.e. we divide a user's UserLimit resources equally among all the apps of the user.

Following image illustrates fair-share assignment for users and apps.



As part of COMPX-4883, we are making following changes inside FifoIntraQueuePreemptionPlugin when FairOrderingPolicy is enabled:

1. In computeAppsIdealAllocation():
 - a. we calculate the FairShare of each app of a user.
2. In skipContainersBasedOnIntraQueuePolicy(), we check that:
 - a. We don't drop an app below its fair-share after preempting its container.