Activeness-based Data Retention Recommender for HPC Facilities

Wei Zhang†
Co-advisors: Yong Chen†, Suren Byna+
Collaborators: Hyogi Sim ✦, Sangkeun Lee ✦, Sudharshan S. Vazhkudai✦
†Texas Tech University. +Lawrence Berkeley National Laboratory. ✦ Oak Ridge National Laboratory.

Motivation

➢ Data retention plays a vital role in retaining valuable files and maintaining sufficient storage space on HPC.

➢ Many data retention solutions have been proposed [2, 8, 9], but have failed to be applied in practice due to significant complexity in their retention criteria or their deployment.

➢ The current in-practice data retention policies [1, 3–5, 7] are variations of the classical time-based data retention solution [6] and hence may cause unnecessary data access interruption or even data loss to users.

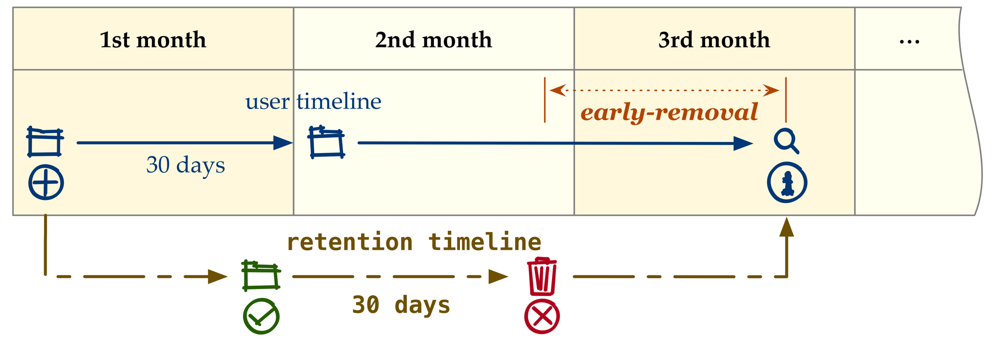

Characteristics

➢ **User Friendliness**

- ActiveDR values active users and their data on the user activeness from multiple perspectives rather than merely some file properties (e.g. timestamps).

➢ **Administrator Friendliness**

- The HPC administrators can easily deploy and run ActiveDR.
- The administrators can selectively execute retention actions from the ActiveDR recommendation result if necessary (e.g. exceptional cases for special users and files).

➢ **Resource Friendliness**

- ActiveDR is the first data retention solution that does not require consideration on files and hence mitigated file scanning operations as well as mitigated metadata server occupation on the file system.
- The activeness evaluation process of ActiveDR runs very fast and does not require much runtime memory.

➢ **Eco Friendliness**

- ActiveDR is the first data retention solution that promotes the active and fruitful use of the HPC facility and hence has significantly positive impact on the HPC ecosystem.

Methods

Overview

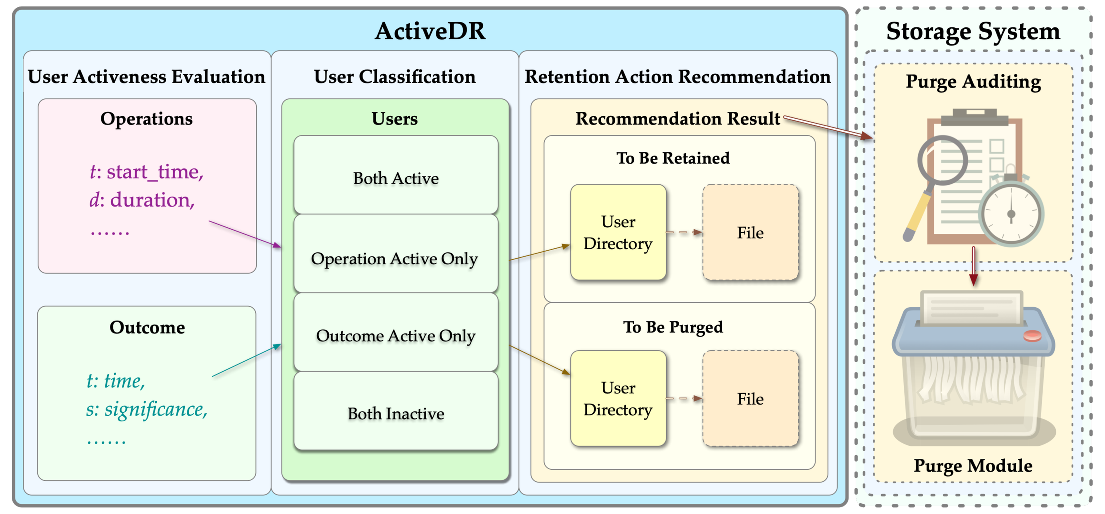

➢ User activeness evaluation framework based on :

❑ the active operations of users on the HPC facility

❑ the outcome delivered by the user operations on the HPC facility

➢ User classification process based on the operation activeness ranking as well as the outcome activeness ranking.

➢ Retention action recommendation process which generates recommendation results about what to be purged and what to be retained.

➢ Software available on Github [10]

User Activeness Evaluation

➢ Essential metrics in user activeness evaluation framework

Meaning	Notation
Set of m Periods	$P = \{p_0, ..., p_{m-1}\}$
Set of n Activity Types	$T = \{\lambda_0, ..., \lambda_{n-1}\}$
Set of k Activities	$A = \{a_0, ..., a_{k-1}\}$
Activeness of Activity a_t at time t	D_{a_t}
Average Activeness of All Activities A in each period	$D_{\overline{A}}$ or $Avg(D_A)$
Activeness Ratio of a Certain Period p	$b_p = D_{A_p}/D_{\overline{A}}$
Period Index of Activity a_t	$e = m - (a_t.ts - a_0.ts)/l + 1$ (ts denotes timestamp)

➢ User activeness rank vector

e	e=5-5+1=1	e=5-4+1=2	e=5-3+1=3	e=5-2+1=4	e=5-1+1=5
b^e	$\left(\frac{D_{tc-5}}{Avg(D_A)}\right)^1$	$\left(\frac{D_{tc-4}}{Avg(D_A)}\right)^2$	$\left(\frac{D_{tc-3}}{Avg(D_A)}\right)^3$	$\left(\frac{D_{tc-2}}{Avg(D_A)}\right)^4$	$\left(\frac{D_{tc-1}}{Avg(D_A)}\right)^5$

$t_c - 5$ $t_c - 4$ $t_c - 3$ $t_c - 2$ $t_c - 1$ t_c

➢ Overall user activeness rank for an activity type λ :

$$\Phi_\lambda = \prod_{e=1}^{m} (b_{p_e})^e$$

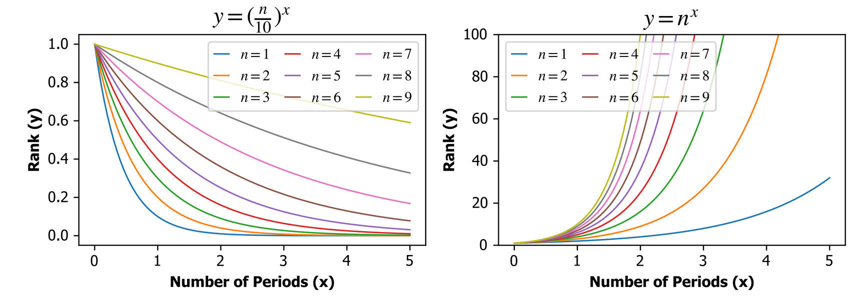

The formula of the user activeness rank is designed based on the monotonic property of exponential function.

- For the activity of type λ, being active contributes more activeness to the overall activeness rank
- For the activity of type λ, the activeness of more recent activities contributes more to the overall activeness rank

User Classification

➢ Calculating operation activeness and outcome activeness

$$\Phi_{op} = \prod_{\lambda_{op}=1}^{m_{op}} \Phi_{\lambda_{op}} \quad and \quad \Phi_{oc} = \prod_{\lambda_{oc}=1}^{m_{oc}} \Phi_{\lambda_{oc}}$$

➢ User classification

➢ Both Active

➢ Operation Active

➢ Outcome Active

➢ Both Inactive

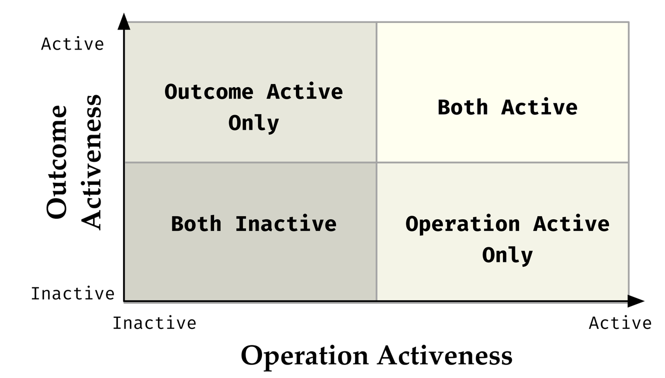

Retention Action Recommendation

➢ Rationale: The more active the user is, the more chance his/her files will survive from being purged.

➢ Activeness-based file lifetime ε_f:

$$\varepsilon_f = \underbrace{d}_{\text{Retention Interval}} \times \Phi_{op} \times \Phi_{oc}$$

➢ Cold-start Problem: Set initial activeness of all activities to 1.0

Evaluation

Experimental Setup

➢ Two-years of user activities (job executions and tracked research publications) from OLCF.

➢ Simulation-based evaluation compared with fixed-lifetime purge policy (a.k.a FLT).

Overall Data Retention Effect

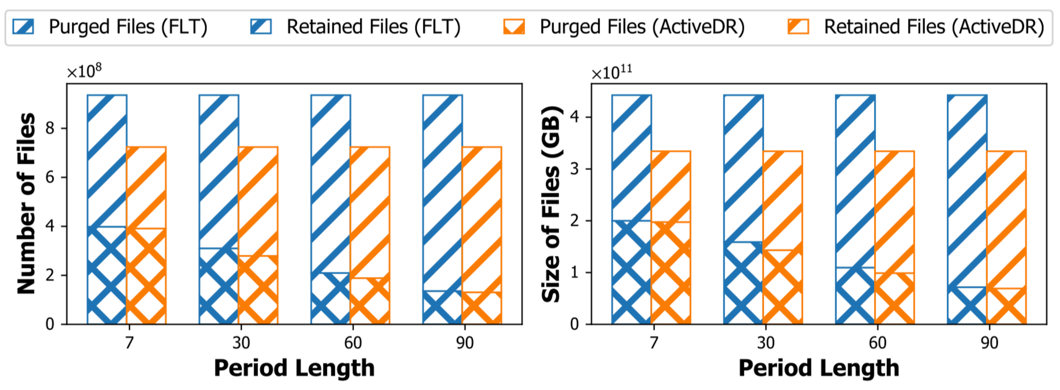

➢ Roughly the same number of files are purged as compared to fixed-lifetime purge policy (a.k.a. FLT)

➢ Less files are scanned for determining the retention actions

Purge Result Break-down

➢ ActiveDR purges more files from inactive users, as compared to FLT.

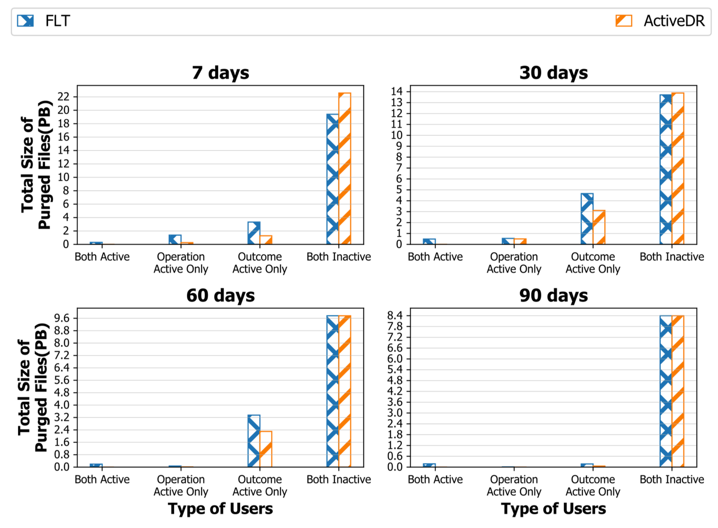

➢ Much fewer active users were affected by file purge operations of ActiveDR, as compared to FLT.

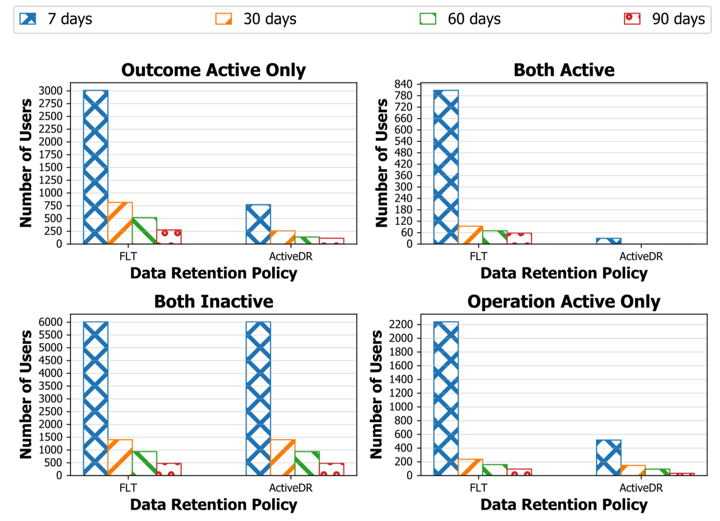

Performance Evaluation

➢ For evaluating 2 years' user activities, ActiveDR only took about 500MB

	User	Publication	Job
Memory (MiB)	48.85	3.5	419.775
Loading Time(s)	4.756075	1.2693	89.877925

➢ Time for loading activities and evaluating the recommendation result takes less than 2 minutes.

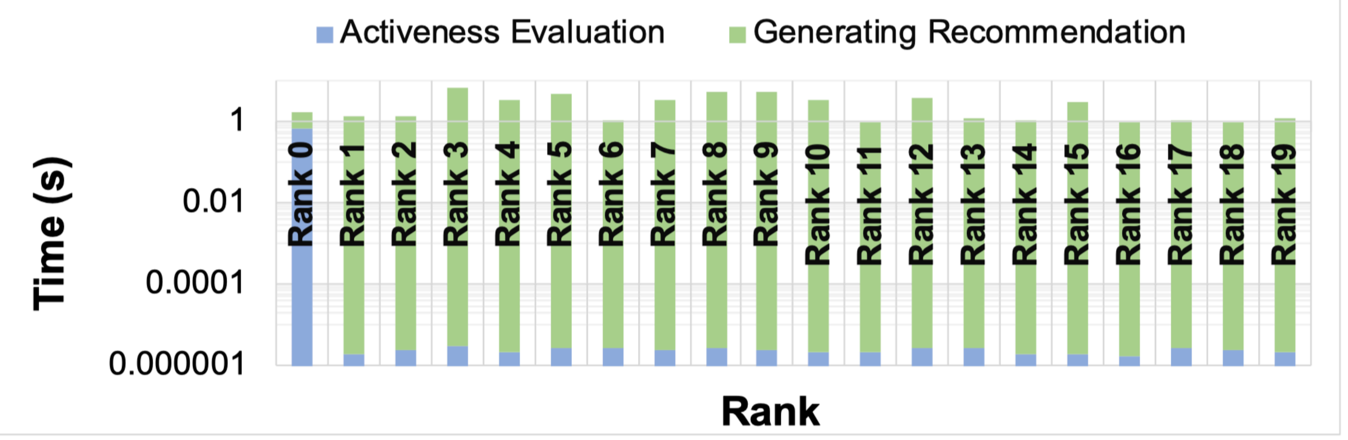

References:

[1] High Performance Computing Center at Texas Tech University. 2020. Data Policy - Use of HPCC Storage. https://www.depts.ttu.edu/hpcc/operations/datapolicy. php.

[2] Henry M. Monti, Ali R. Butt, and Sudharshan S. Vazhkudai. 2009. /Scratch as a Cache: Rethinking HPC Center Scratch Storage. In Proceedings of the 23rd International Conference on Supercomputing (Yorktown Heights, NY, USA) (ICS '09). Association for Computing Machinery, New York, NY, USA, 350–359. https: //doi.org/10.1145/1542275.1542325

[3] National Center for Atmospheric Research. 2020. GLADE File Space. https: //www2.cisl.ucar.edu/resources/storage-and-file-systems/glade-file-spaces.

[4] National Energy Research Scientific Computing Center. 2020. NERSC Data Management Policy. https://docs.nersc.gov/data/policy/.

[5] Oak Ridge Leadership Computing Facility. 2020. OLCF Policy Guides – Data Management Policy.https://docs.olcf.ornl.gov/accounts/olcf_policy_guide.html.

[6] Arie Shoshani, Alexander Sim, and Junmin Gu. 2004. Storage Resource Managers. In Grid Resource Management. Springer, 321–340.

[7] Texas Advanced Computing Center. 2020. TACC Usage Policy. https://portal. tacc.utexas.edu/tacc-usage-policy.

[8] Lars Arne Turczyk, Oliver Heckmann, and Ralf Steinmetz. 2007. File Valuation in Information Lifecycle Management. In Proceedings of the Thirteenth Americas Conference on Information Systems, Keystone, Colorado.

[9] Fons Wijnhoven, Chintan Amrit, and Pim Dietz. 2014. Value-Based File Retention: File Attributes as File Value and Information Waste Indicators. J. Data and Information Quality 4, 4, Article 15 (May 2014), 17 pages. https://doi.org/10.1145/2567656

[10] **Wei Zhang. 2020. ActiveDR (v1.0.2). https://doi.org/10.5281/zenodo.3878756**

Acknowledgement:

We are thankful to the anonymous reviewers for their valuable feedback. This research is supported in part by the National Science Foundation under grant CCF-1409946, CCF-1718336, and CNS-1817094.

This work is supported in part by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC0205CH11231..

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility.