

D3N: A multi-layer cache for improving big-data applications' performance in data centers with imbalanced networks

Matt Benjamin, Red Hat

Emine Ugur Kaynar, Boston University

Ali Maredia, Red Hat

Abstract

Caching methods for improving the performance of big-data jobs when accessing data lakes assume unlimited bandwidth across the datacenter to cache nodes. However, many enterprise and academic datacenters exhibit significant network imbalances due to varying load, oversubscribed networks, and uncoordinated networking upgrades. This work describes D3N, a multi-layer caching architecture that caches data on the access side of potential network bottlenecks in hierarchical datacenter networks. Co-locating per-layer caches allows individual cache sizes to be dynamically tuned based on observed conditions. To understand the benefits of multi-layer caching, the paper describes and evaluates an implementation of a two-layer version of D3N within a commonly used Datalake, Ceph. It also describes and evaluates an algorithm for dynamically tuning a co-located two-layer cache.

Speaker Bios

Matt Benjamin is the engineering manager for the Ceph Object Storage and user-space NFS teams at Red Hat, and past PI for CohortFS, an NFS-funded research project that explored parallel NFS access and performance optimizations in the Ceph distributed storage system.

Emine Ugur Kaynar is a PhD student in Computer Science at Boston University. She is a member of the Systems Research Group working with Prof. Orran Krieger as her advisor. Her research interests lie broadly in the fields of cloud computing, distributed systems and data-intensive computing. Her current research focuses on i) improving storage efficiency in a data center for bigdata analytic workloads by designing storage cache architectures, ii) exploring the performance characteristic of erasure coded storage systems. Before joining Boston University, She received her MSc from State University of New York at Binghamton, Department of Computer Science advised by Prof. Ping Yang and my BSc from Bogazici University, Department of Information Systems Engineering.

Ali Maredia is a Software Engineer at Red Hat working on Ceph and Object Storage. He has been working on Ceph distributed storage system for 4 years. Before Red Hat he was at CohortFS, a NFS-funded project that explored parallel NFS access and performance optimizations in Ceph.