



Comparing Cache Utilization Trends for Regional Data Caches

John Wu

A. Sim, E. Wang, R. Monga, J. Balcas, C. Guok, I. Monga, D. Davila, F. Wurthwein, H. Newman

CHEP 2024

Introduction

- **Computational centers and experimental facilities are inundated with data**
 - Data acquisition, storage, and analysis is becoming increasingly important for DOE science missions, as data volumes continue to grow
 - Energy Science network (ESnet) aims to support advanced distributed data management services beyond simply providing bandwidth
 - To place the data close to users with in-network caching service (focus of this work)
 - To explore options for predictive storage (future)
- **Goals of this caching pilot**
 - Study the effectiveness of caches in reducing wide-area network traffic
 - Characterize in-network caching data access patterns
 - Help plan for additional deployments of caches in the science network infrastructure
 - Exercise cache management software: DTN-as-a-Service (DTNasS).
 - Explore the predictability of resource utilization to design predictive caching policies
- **Observations**
 - Significant portions of popular datasets are re-used in the user analyses
 - Storage cache allows data sharing among users in the same geographical region
 - Reduce the traffic on finite capacity of the wide-area network
 - Improve overall application performance by decreasing access latency and increasing analysis throughput

Impact of In-network Caching Pilot

- **Collaborating with HEP communities to improve their distributed workflows**
- **Validated the benefits of regional caching for a scientific workload**
 - **Initial collaboration with US CMS in Southern California**
 - **Implement a rigorous approach to cache utilization analysis**
 - **Recommended additional cache storage that improved cache hit rates to 94% based on data volume**
 - **Results prompted the additional deployments of caching nodes in Boston and Chicago, with additional expansion for Amsterdam and London for LIGO and DUNE experiments through OSDF**
 - **Similar cache utilization studies are planned to validate the feasibility of transitioning the pilot into a production service offering**



Current In-network Cache Installations

- **US CMS: Southern California Petabyte Scale Cache**
 - Regional storage cache for US CMS user analysis at Caltech and UCSD
 - 23 federated XCache nodes: Approximately 2PB of total storage capacity
- **US CMS: Chicago**
 - Regional cache for U. Wisc Madison, Purdue and Notre Dame
 - 6 federated XCache nodes: Approximately 345 TB in total
- **US CMS: Boston**
 - Regional cache for MIT
 - 1 XCache node: 170TB on LHCONE
- **US ATLAS: Boston (still testing, no access logs yet)**
 - Regional cache for BNL
 - 1 XCache node: 170TB on LHCONE



Data Access Summary Statistics for US CMS


		# of accesses	Cache misses (TB)	Cache hits (TB)	% of cache hits
6/2020-3/2024 SoCal Cache	Total	30,938,524	11,837.38	26,686.73	69.3%
	Daily average	22,114	8.46	19.08	
10/2022-3/2024 Chicago Cache	Total	9,054,942	13,315.37	12,494.41	48.4%
	Daily average	16,553	24.34	22.84	
8/2023-3/2024 Boston Cache	Total	27,754,028	87,311.58	6,173.02	6.6%
	Daily average	114,214	359.31	25.40	

US CMS - SoCal Cache (6/2020-3/2024)




SoCal: daily data access requests


Up to 100K requests per day on xcache-11



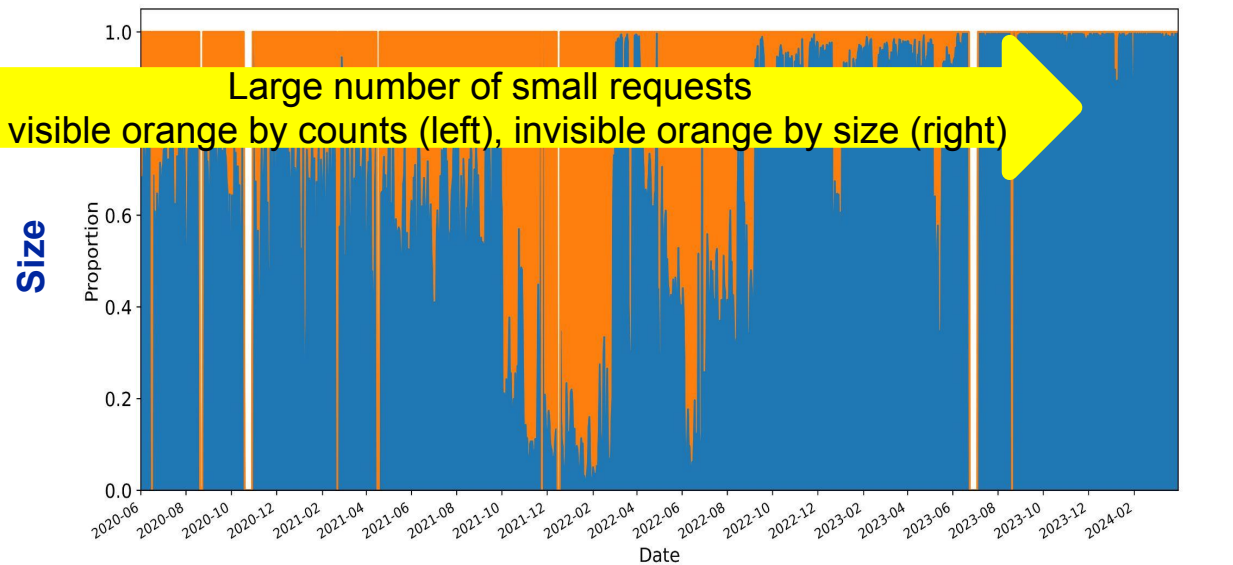



Up to 250 TB requested per day on xcache-11



Average data access size between 1 and 4 GB



SoCal: daily cache hits and cache misses



Large number of small requests
-- visible orange by counts (left), invisible orange by size (right)

US CMS - Chicago Cache (10/2022-03/2024)

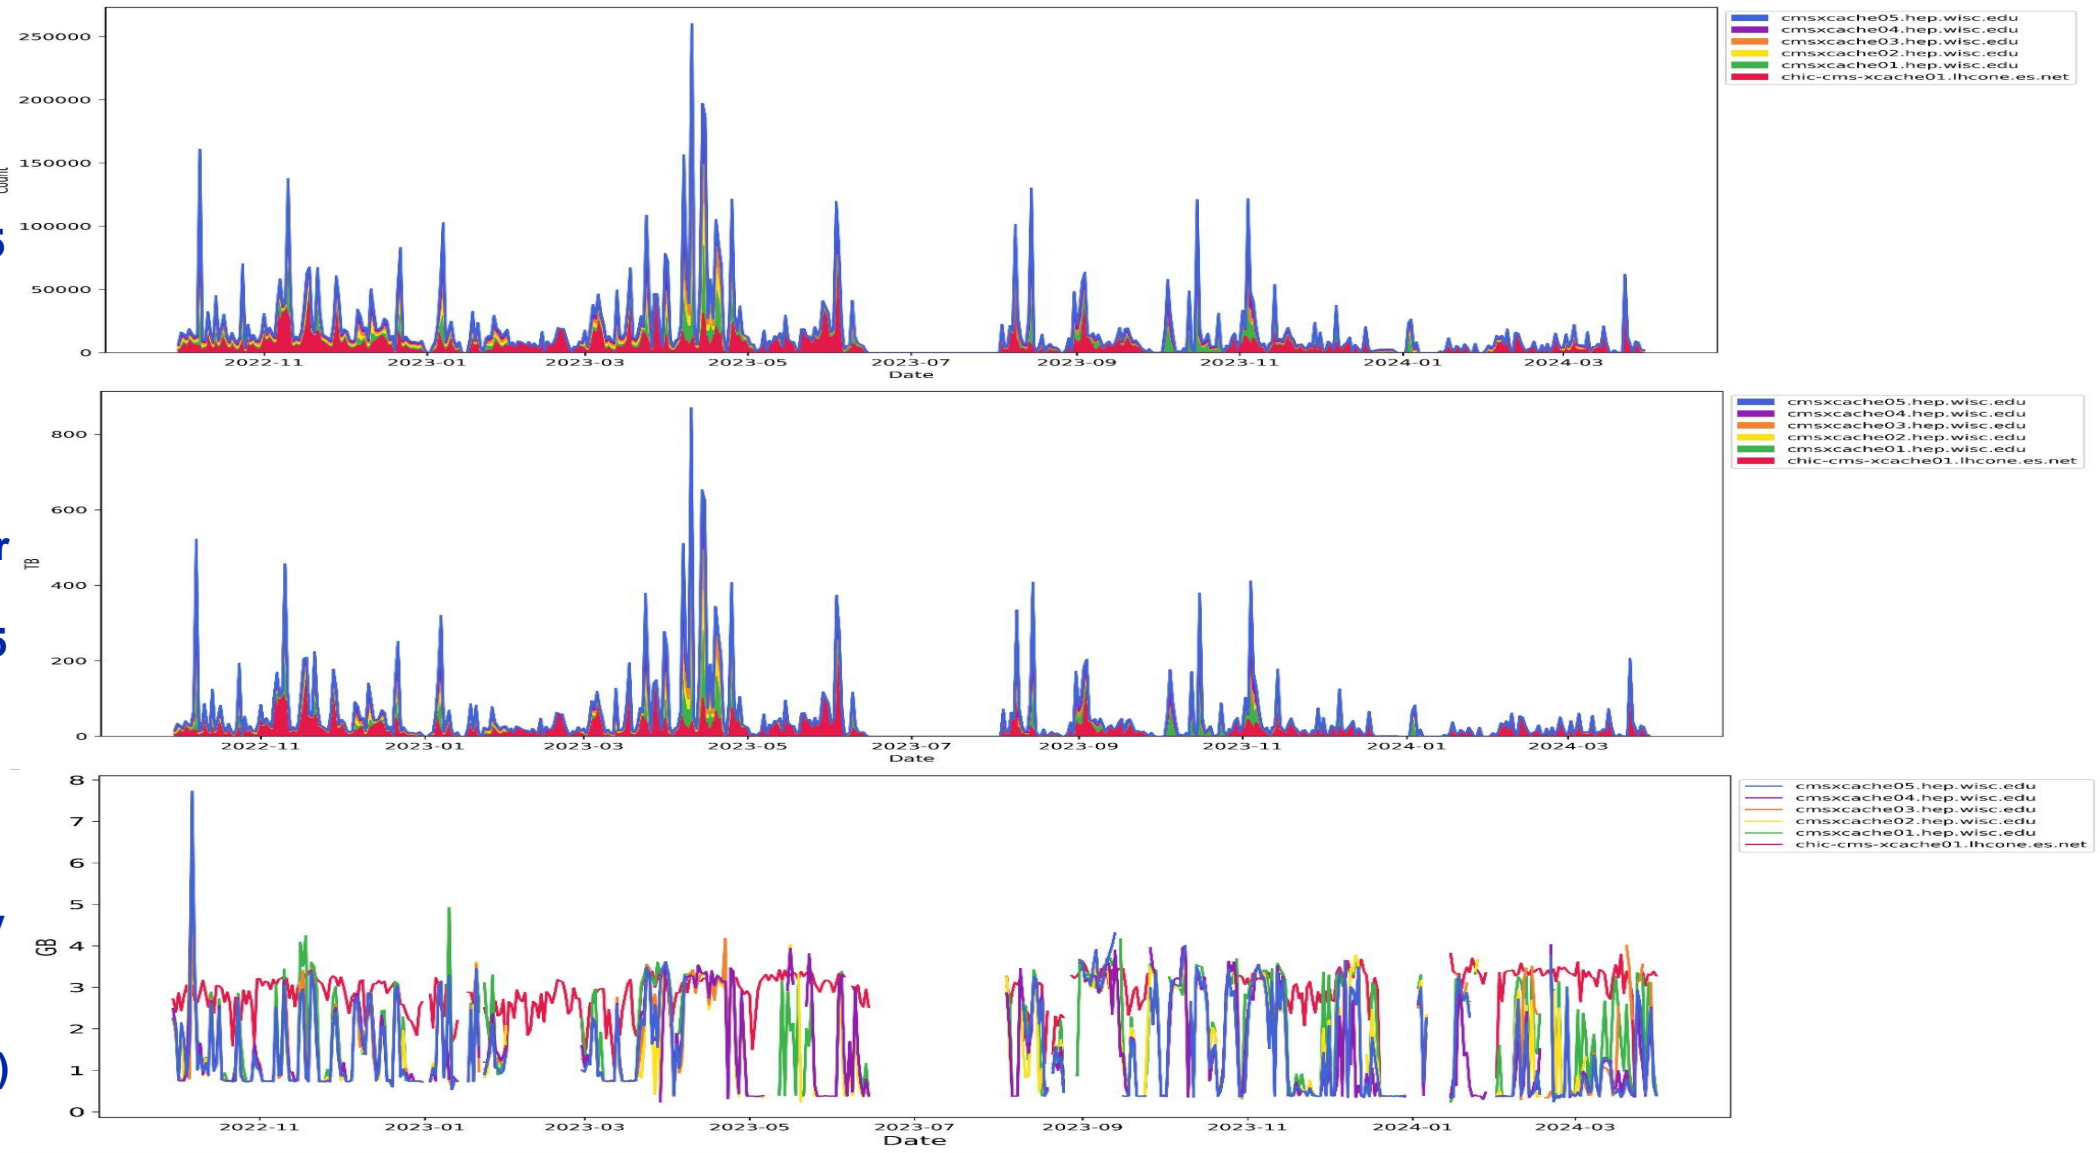


Chicago: daily data access requests

Up to 250K requests per day on cmsxcache05

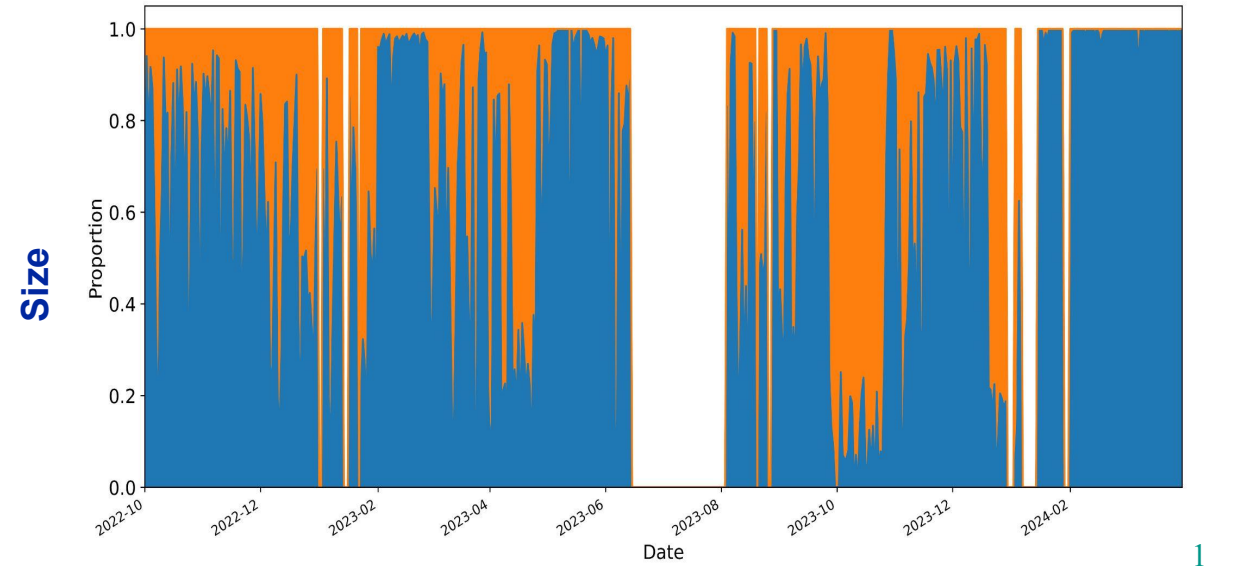

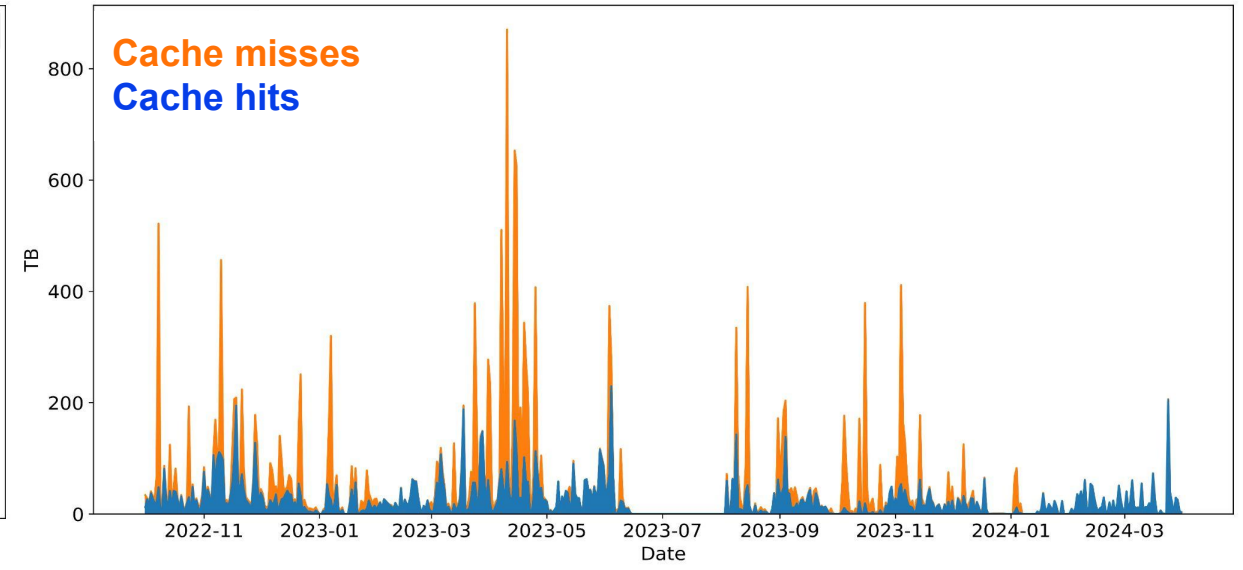

Up to 800 TB requested per day on cmsxcache05

Average request size occasionally goes above 4GB (unexpected)





Chicago: daily cache hits and cache misses




US CMS - Boston Cache (8/2023-03/2024)




Boston: daily data access requests


Up to 600K requests per day



Up to 2000 TB requested per day





average size between 3 and 4 GB

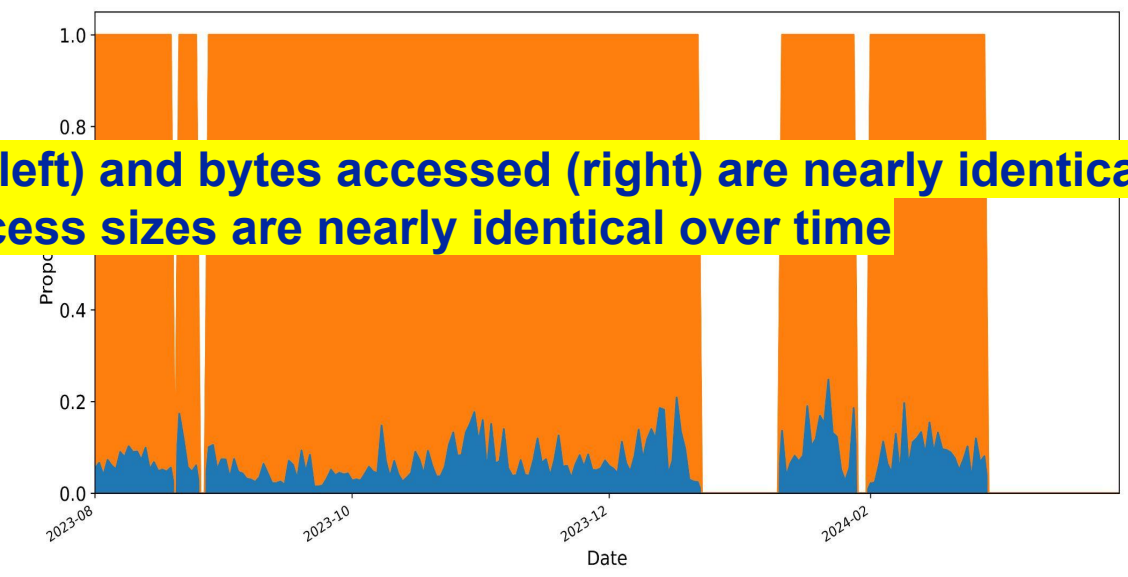
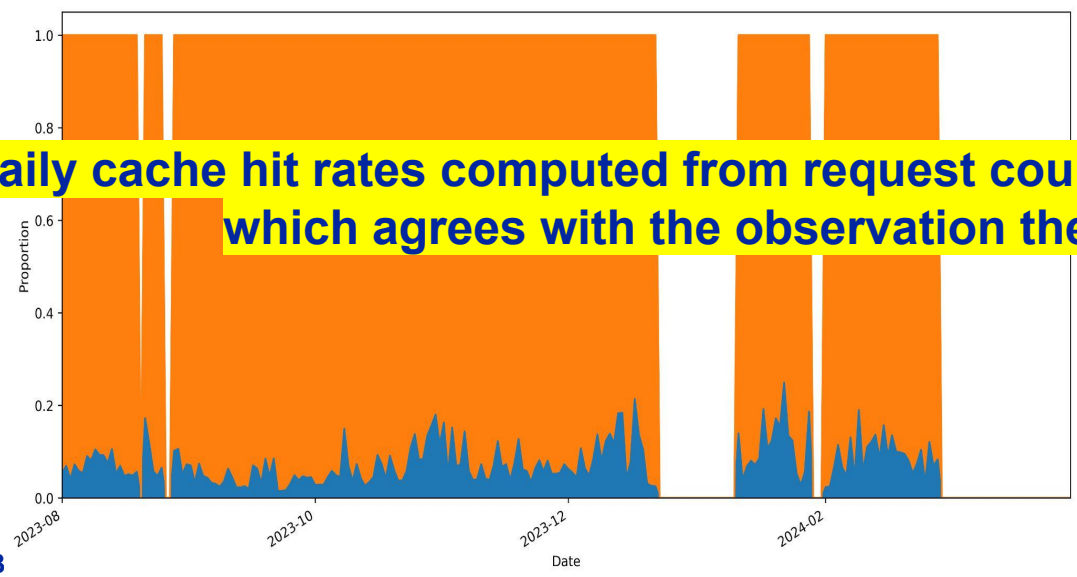




Boston: hourly cache hit rates computed from request counts and bytes accessed nearly identical



Daily cache hit rates computed from request counts (left) and bytes accessed (right) are nearly identical which agrees with the observation the access sizes are nearly identical over time




Characteristics of Network and Cache Utilization

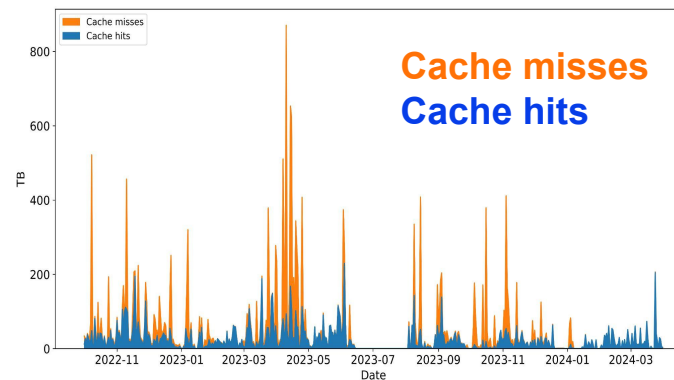
-- Comparison of three caching sites



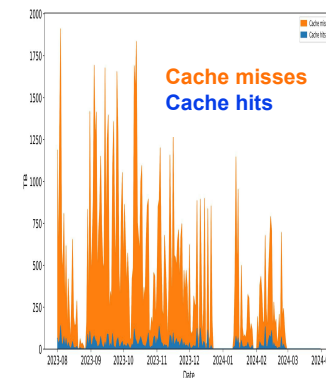
Daily Volume of Cache Hits and Misses



SoCal Cache from June 2020 to Mar. 2024, with 30.94 million data access records for 11.8PB of traffic volume for cache misses and 26.7PB of traffic volume for cache hits. 69.3% of the total traffic has been saved from the cache of 2PB.




Chicago Cache from Oct. 2022 to Mar. 2024, with 9.05 million data access records for 13.3PB of traffic volume for cache misses and 12.5PB of traffic volume for cache hits. 48.4% of the total traffic has been saved from the cache of 340TB.



Boston Cache from Aug. 2023 to Mar. 2024, with 27.75 million data access records for 87.3PB of traffic volume for cache misses and 6.2PB of traffic volume for cache hits. 6.6% of the total traffic has been saved from the cache of 170TB.




Ratio of Daily Volume of Cache Misses and Hits




SoCal Cache from June 2020 to Mar. 2024, with 30.94 million data access records for 11.8PB of traffic volume for cache misses and 26.7PB of traffic volume for cache hits. 69.3% of the total traffic has been saved from the cache of 2PB.

Cache misses
Cache hits



Chicago Cache from Oct. 2022 to Mar. 2024, with 9.05 million data access records for 13.3PB of traffic volume for cache misses and 12.5PB of traffic volume for cache hits. 48.4% of the total traffic has been saved from the cache of 340TB.



Boston Cache from Aug. 2023 to Mar. 2024, with 27.75 million data access records for 87.3PB of traffic volume for cache misses and 6.2PB of traffic volume for cache hits. 6.6% of the total traffic has been saved from the cache of 170TB.

Predictability of Network and Cache Utilization

SoCal Cache

Prediction models --



- LSTM tensorflow/keras implementation
- Daily prediction: 128 units, with tanh activation function, 0.04 dropout rates, and the number of epochs of 50.
- Hourly prediction: 64 units, with relu activation function, 0.1 dropout rates, and the number of epochs of 50.





SoCal daily request prediction

Blue: measurements; Orange: training; Red: prediction



Total accesses



Cache Hits



Cache Misses









Chicago daily request prediction

Blue: measurements; Orange: training; Red: prediction



Total accesses



Cache Hits



Cache Misses









Boston daily request prediction

Blue: measurements; Orange: training; Red: prediction



Total accesses



Cache Hits



Cache Misses



Summary and Future Plans

- From the SoCal Cache deployment, it reduced about 33 TB of network traffic on average from wide-area network everyday during the last 11 months of the study period.
 - From the network bandwidth point of view, this number is not that high.
 - Translates to about 3Gbps network capacity saved on average.
 - Peak bandwidth might be a better measure.
 - However, it has a significant impact to the science community because 2/3rd of the file requested could be served from a nearby data cache. After increasing the cache capacity, nearly 94% of the accesses are cache hits.
- The success of SoCal Cache is generating interests from other science applications for data caching
 - Predictive data placement could make the service even more useful
- Generalize the data service software for the science community.