Network and Cache Utilization Trends of Regional Scientific Data Cache for US-CMS SoCal HEP Analysis Jobs

Caitlin Sim
University of California, Berkeley

Advisors and collaborators
ESnet: Chin Guok, Inder Monga
LBNL: Alex Sim, John Wu
UCSD: Frank Würthwein, Diego Davila

July 22, 2022

Introduction

- **Large data volume generation in scientific experiments and simulations**
  - Challenging for geographically distributed collaborations
    - E.g., Large Hadron Collider (LHC) from High-Energy Physics (HEP) community
  - Data stored at a few locations
    - Requiring significant networking resources
    - ATLAS Tier-1 site at Brookhaven National Laboratory, USA
    - CMS Tier-1 site at Fermi National Accelerator Laboratory, USA
    - Network traffic primarily carried by Energy Sciences Network (ESnet)
- **Observation**
  - Significant portion of the popular dataset is transferred multiple times
  - Storage cache allows data sharing between users in same region
    - Reduce the redundant data transfers over the wide-area network
    - Save network traffic volume
    - Lower data access latency
    - Improve overall application performance
  - Use case: Southern California Petabyte Scale Cache (SoCal Repo)

Introduction (2)

- **Goals**
  - Explore measurements from Southern California Petabyte Scale Cache (SoCal Repo) to understand
  - Trends of network utilization
  - Effectiveness of the SoCal Repo in reducing network traffic
  - Predictability of traffic patterns for uses beyond HEP community
- **Measurement data**
  - 1 year logs from SoCal Repo nodes from July 2021 - June 2022
  - Analysis on Cori and Perlmutter at NERSC

Southern California Petabyte Scale Cache (SoCal Repo)

- **High-Luminosity Large Hadron Collider**
  - HL-LHC aims to increase performance after 2025
- **SoCal Repo consists of 24 federated storage caches for US CMS**
  - 12 nodes at UCSD: each with 24 TB, 10 Gbps network connection
  - 11 nodes at Caltech: each with storage sizes ranging from 96TB to 388TB, 40 Gbps network connections
  - 1 node at ESnet: 44 TB storage, 40 Gbps network connection
  - Approximately 2.5PB of total storage capacity
  - ~100 miles between UCSD and Caltech nodes, round trip time (RTT) < 3 ms
  - ~460 miles between ESnet and UCSD nodes, RTT ~10 ms

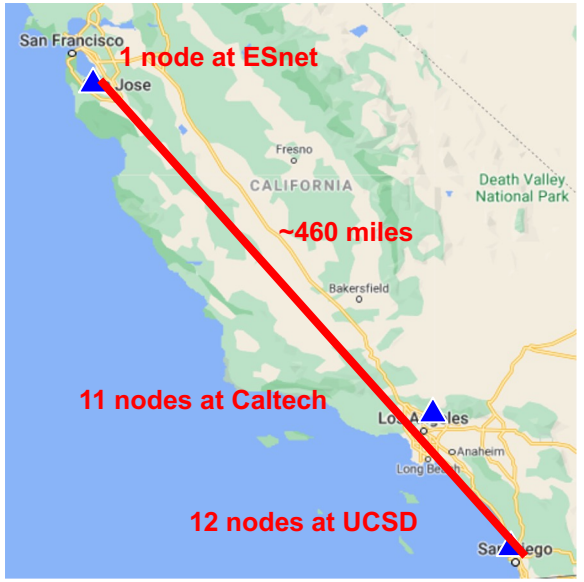

Sunnyvale–San Diego is the relevant distance scale

- **Working with US CMS data analysis using MINIAOD/NANOAOD**
  - One Caltech node is for NANOAOD, and the rest are for MINIAOD
  - Analysis Object Data (AOD):
    - 384 PB of RAW
    - 240 PB of AOD

    } Mostly on Tape: accessed a few times per year
    - 30 PB of MINIAOD
    - 2.4 PB of NANOAOD

    } Mostly on disk: heavily re-used by many researchers

Background in data caching

- **Diverse science relevant to DOE HEP & NP**
  - Regional storage repo and data caching are one of the hot topics at HSF/WLCG meetings
  - At present, there are caches in production for ATLAS, CMS, and OSG, all based on XRootD
    - OSG cache use dominated by Dune, LIGO, Virgo, MINERVA, DES, NOVA, and a liquid XENON detector R&D for future dark matter and neutrinoless double beta decay experiments. Electron Ion Collider R&D is in planning

Collaboration	Working Set	Data Read	Reread Multiplier
DUNE	25 GB	131 TB	5400
LIGO (private)	41.4 TB	3.8 PB	95
LIGO (public)	4.3 TB	1.5 PB	318
MINERVA	351 GB	116 TB	340
DES	268 GB	17 TB	66
NOVA	268 GB	308 TB	1200
RPI_Brown	67 GB	541 TB	8300

OSG Data Federation Caches
are deployed at network POPs (yellow)
and compute endpoints (blue).

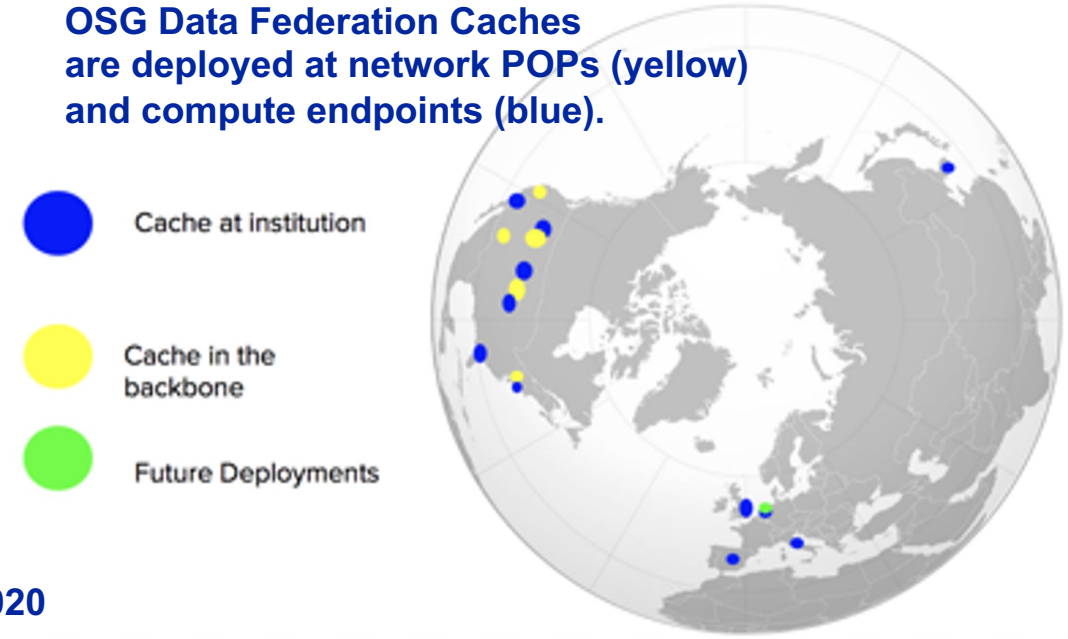

Data read from OSG data federation caches in 6 month period 3/2020-8/2020

Cache hits and cache misses

- **Cache hits**
  - Shared data access
  - Repeated data accesses
  - No need to transfer from remote sites
  - Data traffic from the local node cache to the application
- **Cache misses**
  - First time data access
  - Data traffic from remote sites to the local node cache, then to application

Summary data accesses July 2021 - June 2022

	Number of accesses	Data transfer size (TB)	Shared data size (TB)	Percentage of shared data size
July 2021	1,182,717	385.78	519.25	57.37%
Aug 2021	1,078,340	206.94	313.46	60.23%
Sept 2021	1,089,292	206.96	257.18	55.41%
Oct 2021	1,058,071	412.18	141.91	25.61%
Nov 2021	878,703	649.30	82.67	11.29%
Dec 2021	983,723	1,257.89	130.03	9.37%
Jan 2022	1,207,332	2,238.59	148.26	6.21%
Feb 2022	451,495	320.65	21.40	6.25%
Mar 2022	370,228	64.07	444.65	87.41%
Apr 2022	1,469,630	596.74	1,412.96	70.31%
May 2022	688,543	877.23	865.07	49.65%
June 2022	415,022	994.42	162.57	14.05%
Total	8,021,922	8,210.78	4,499.44	35.40%
Daily average	22,283.12	22.81	12.46	

Daily number of data accesses

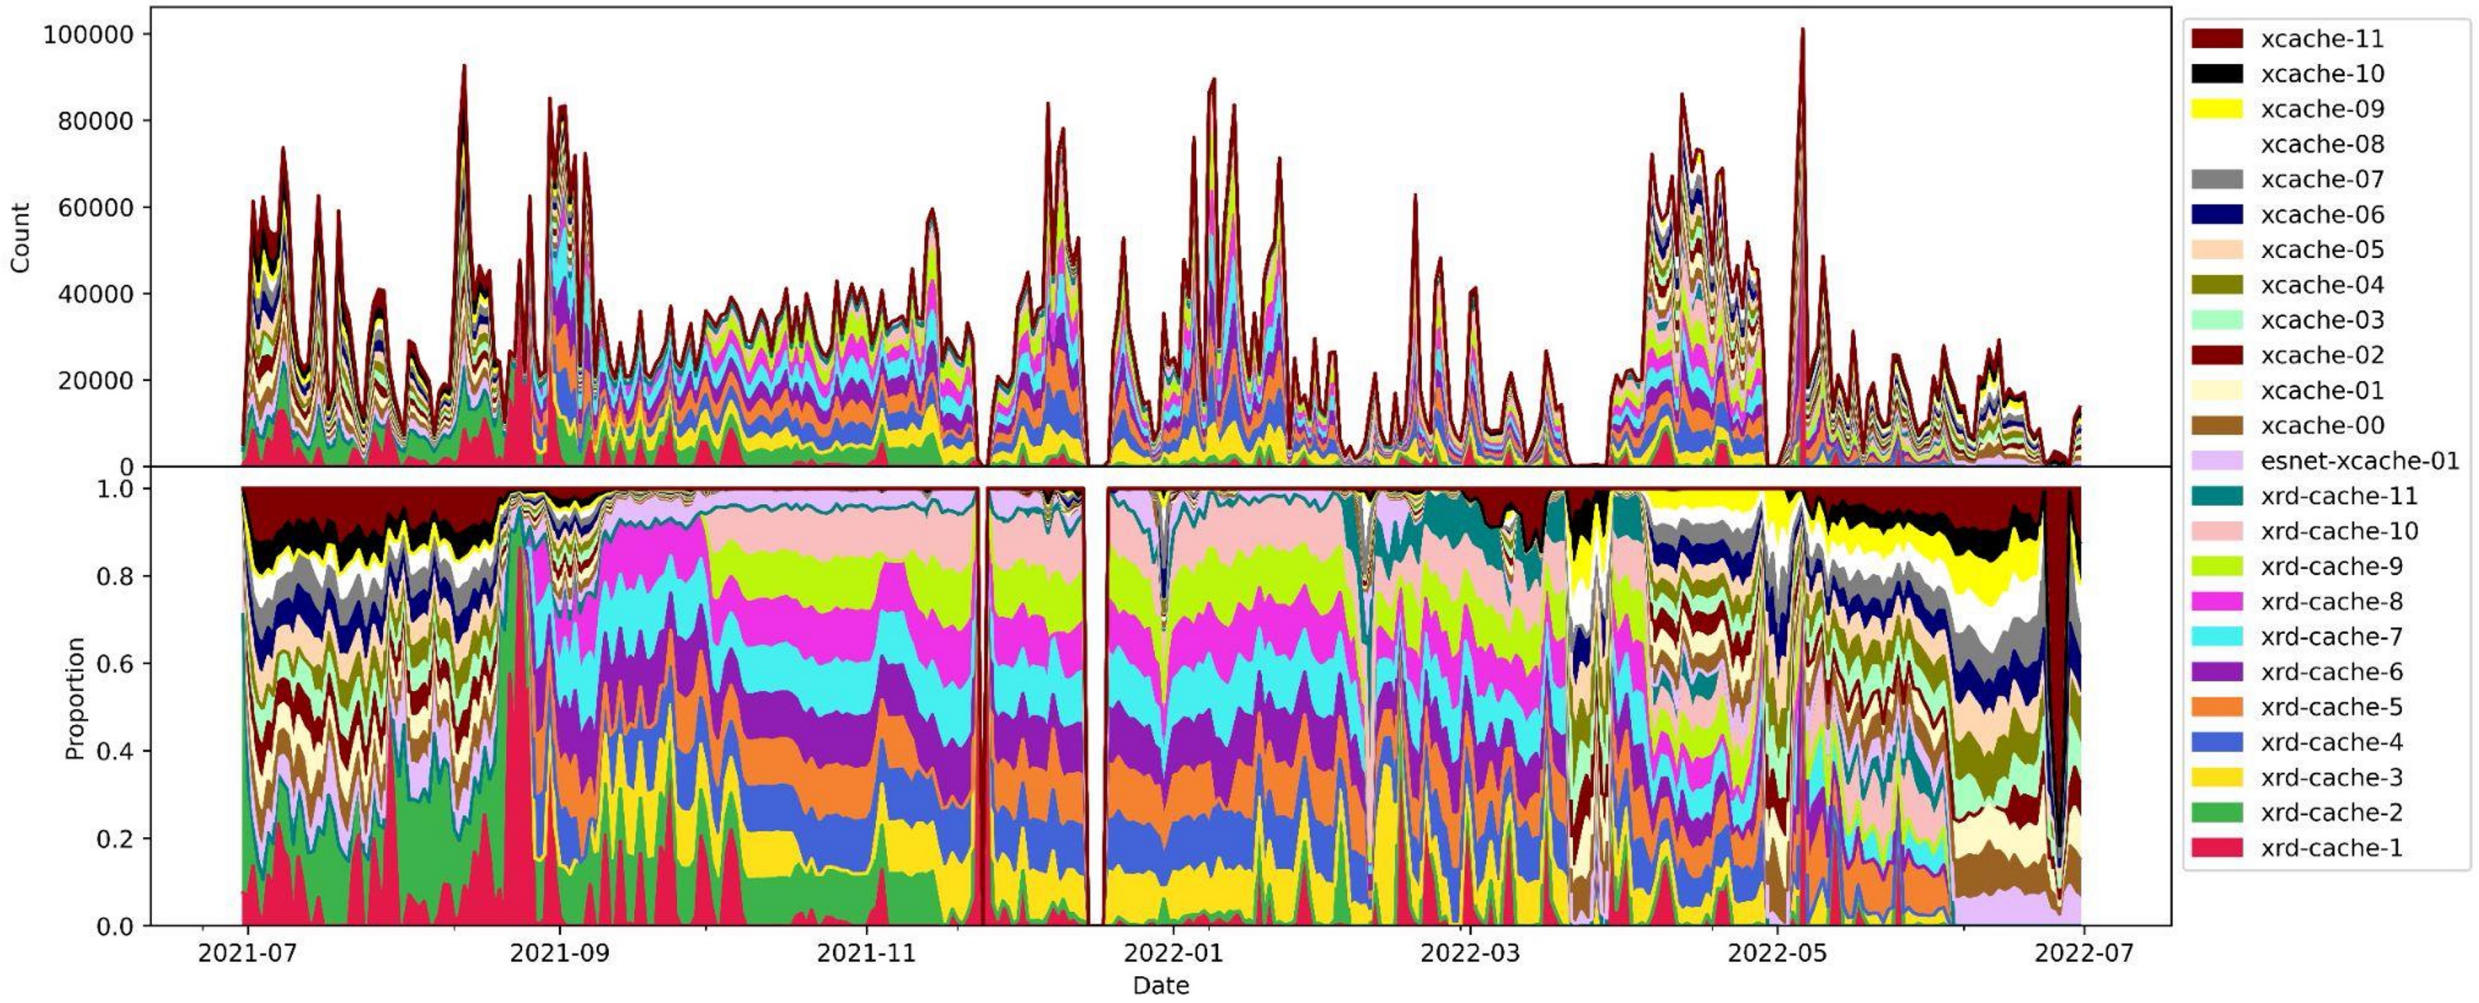

Weekly number of data accesses

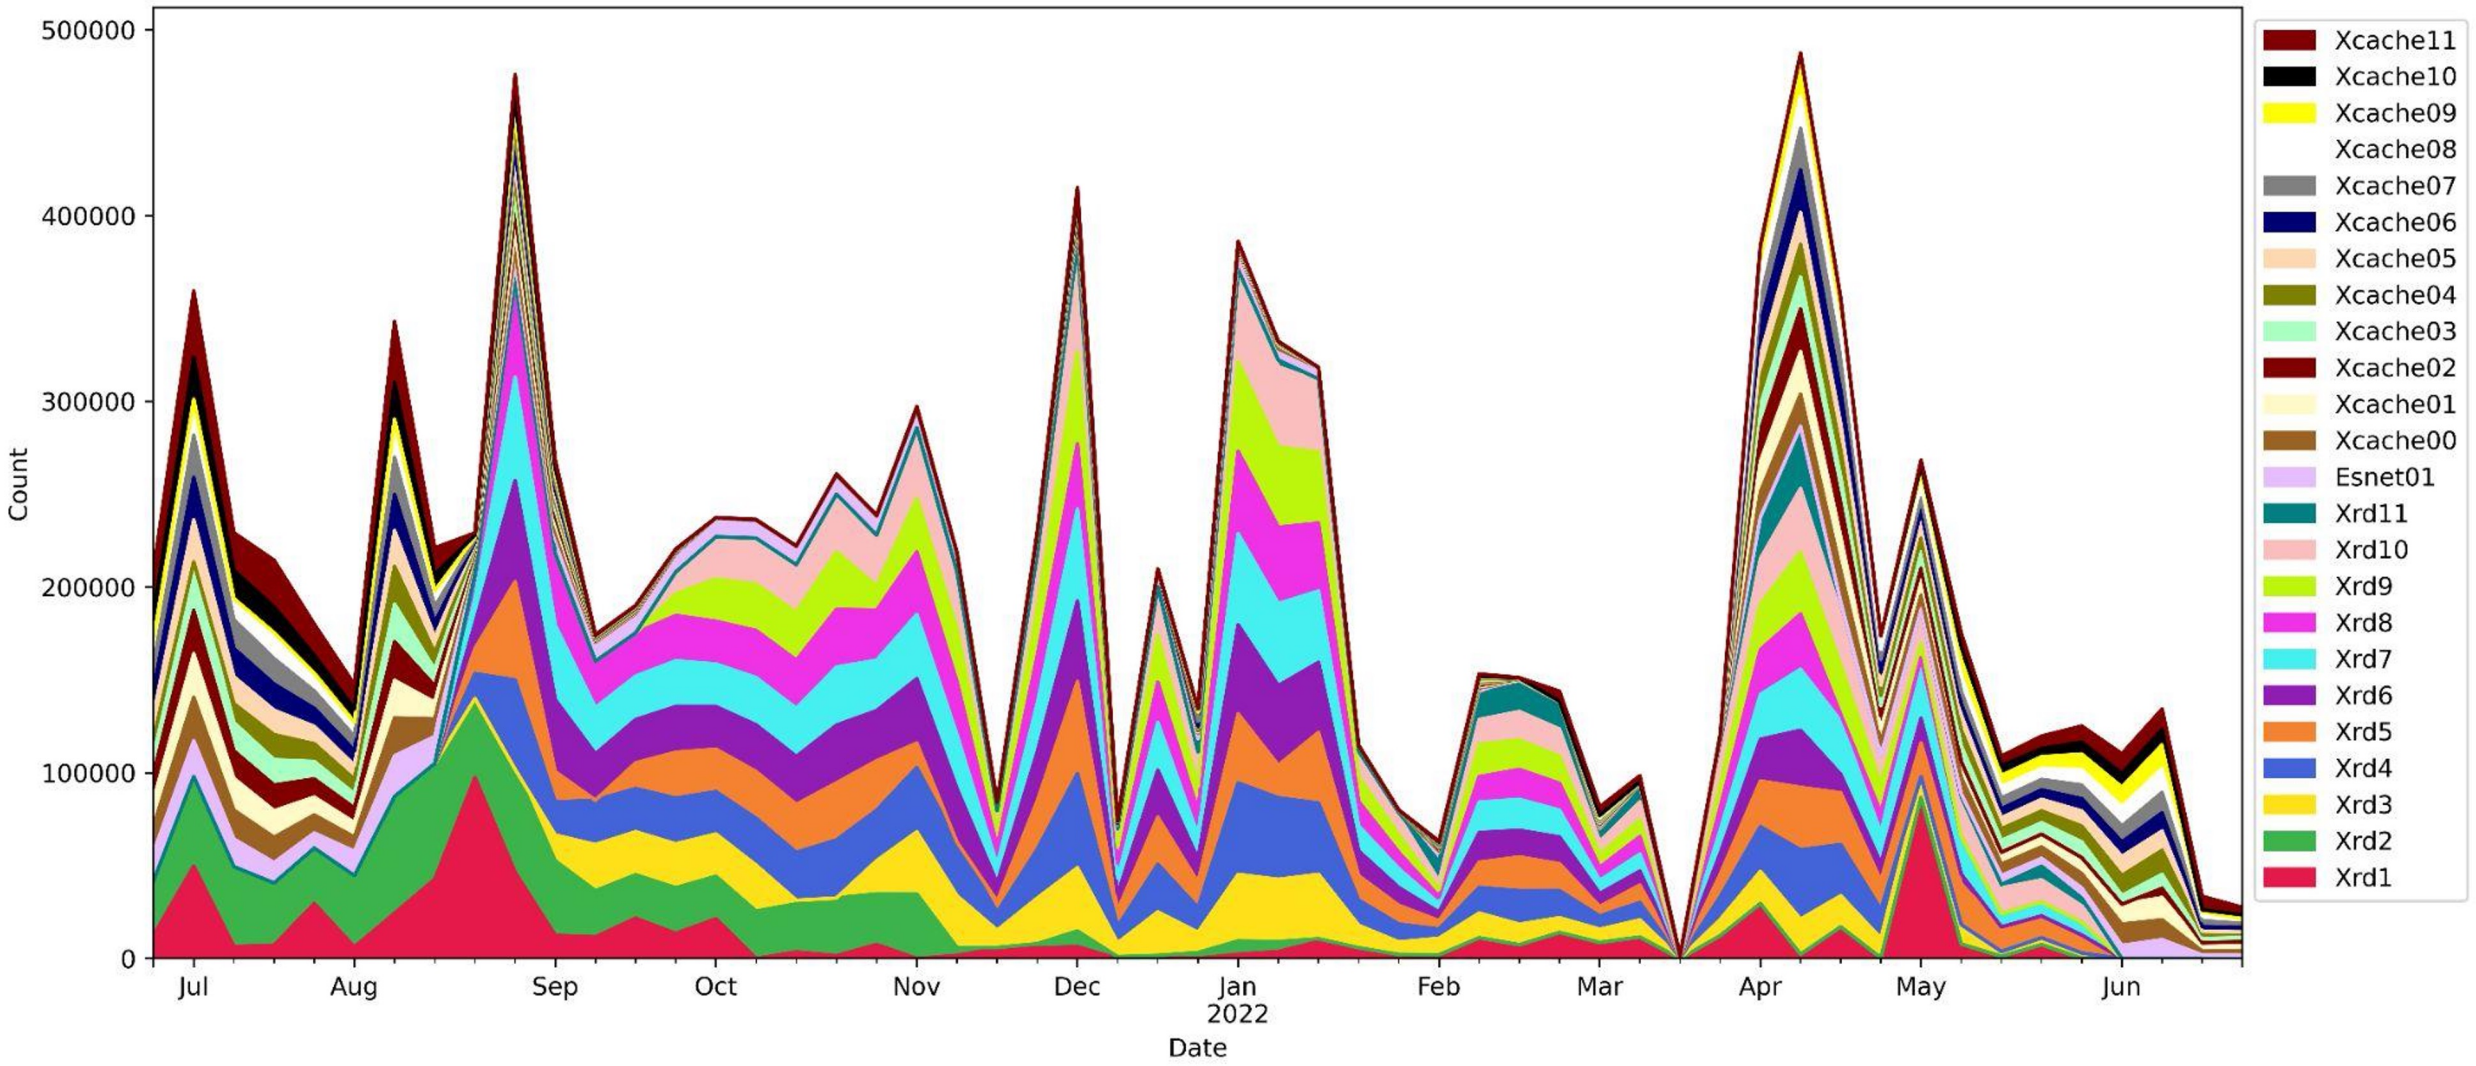

Daily data access volume

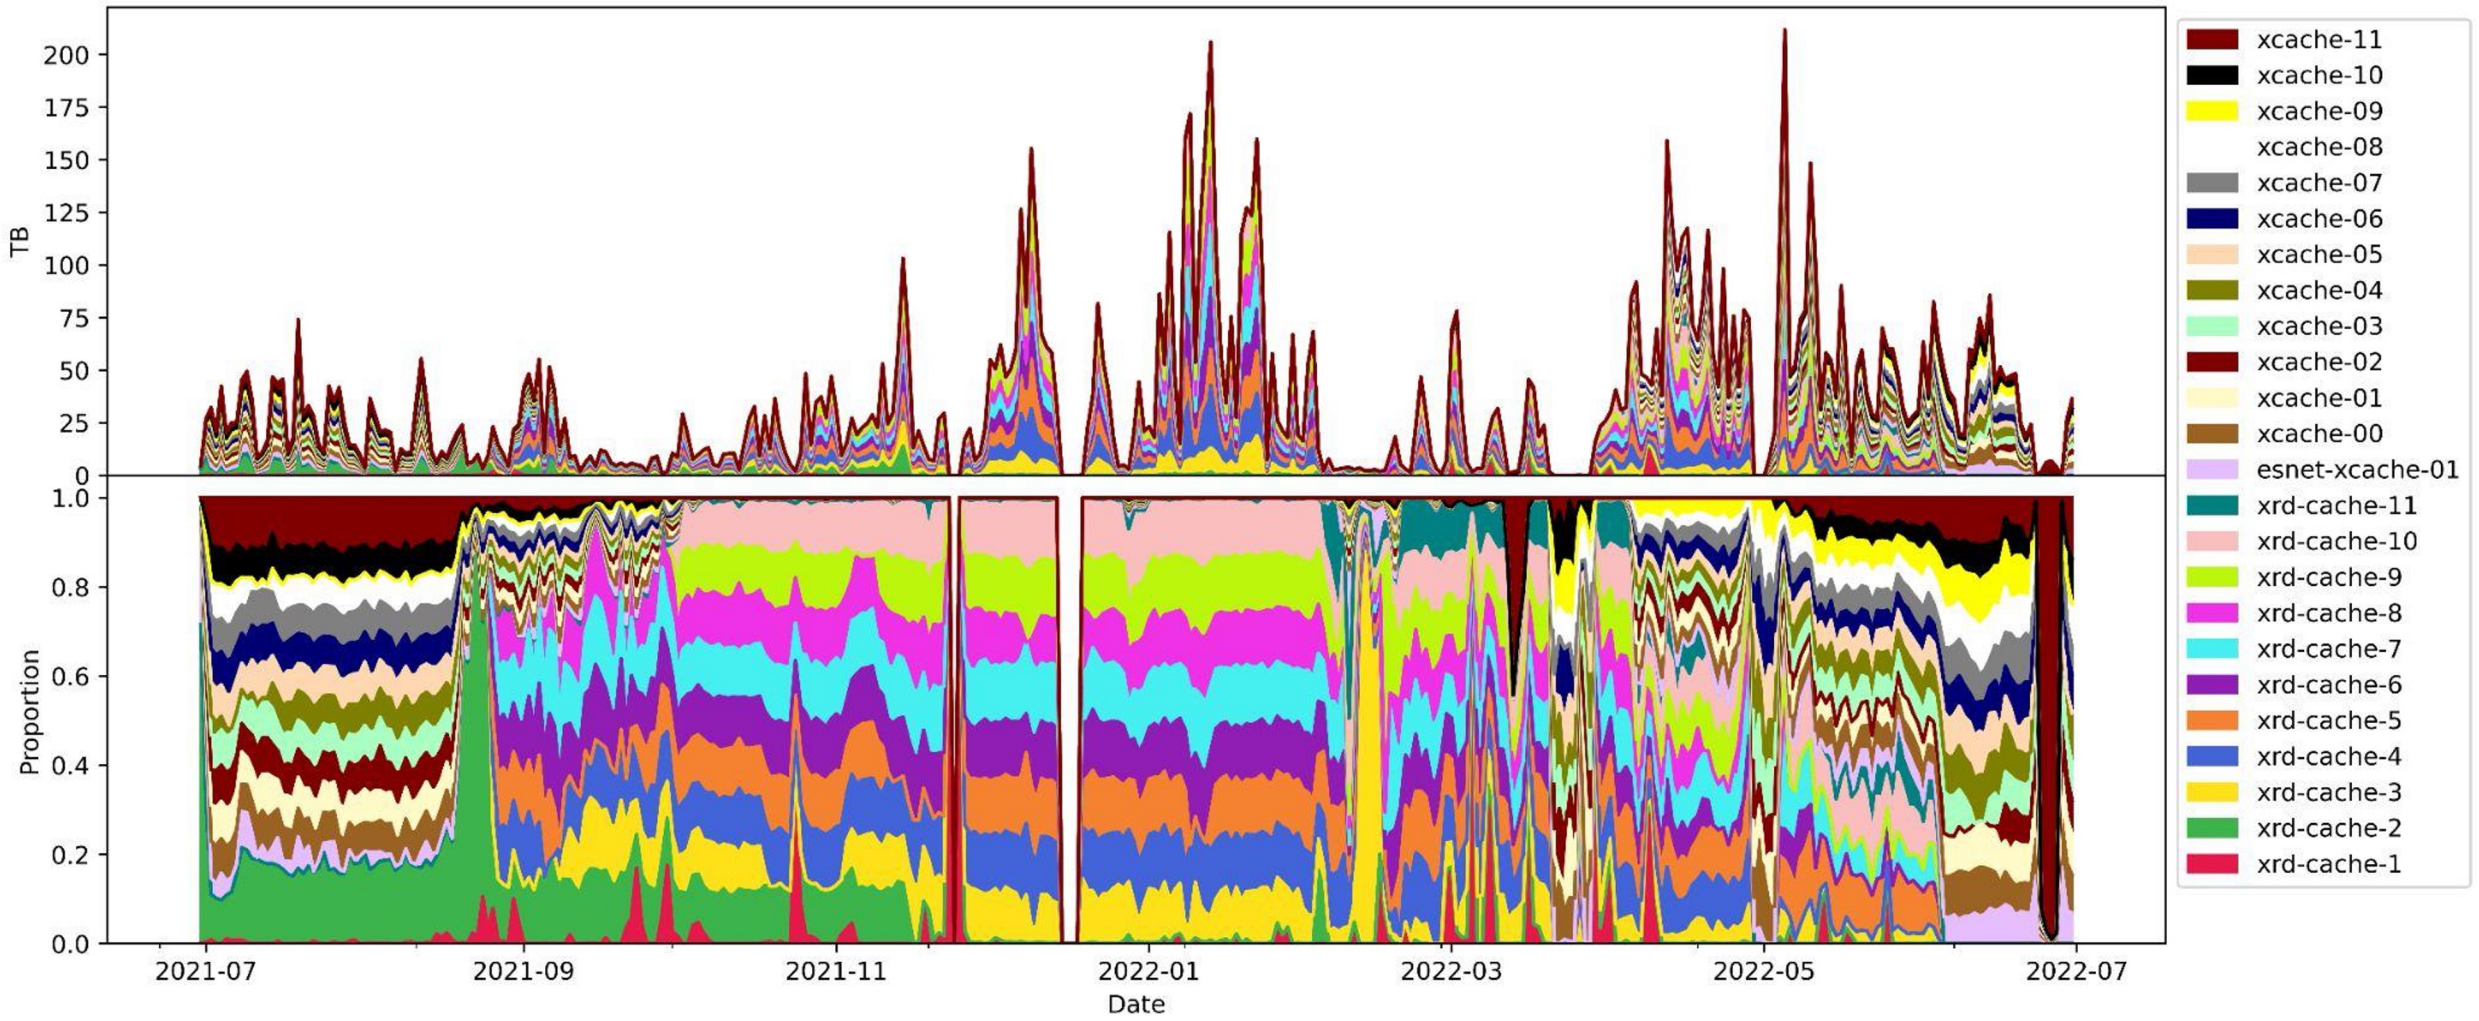

Weekly data access volume

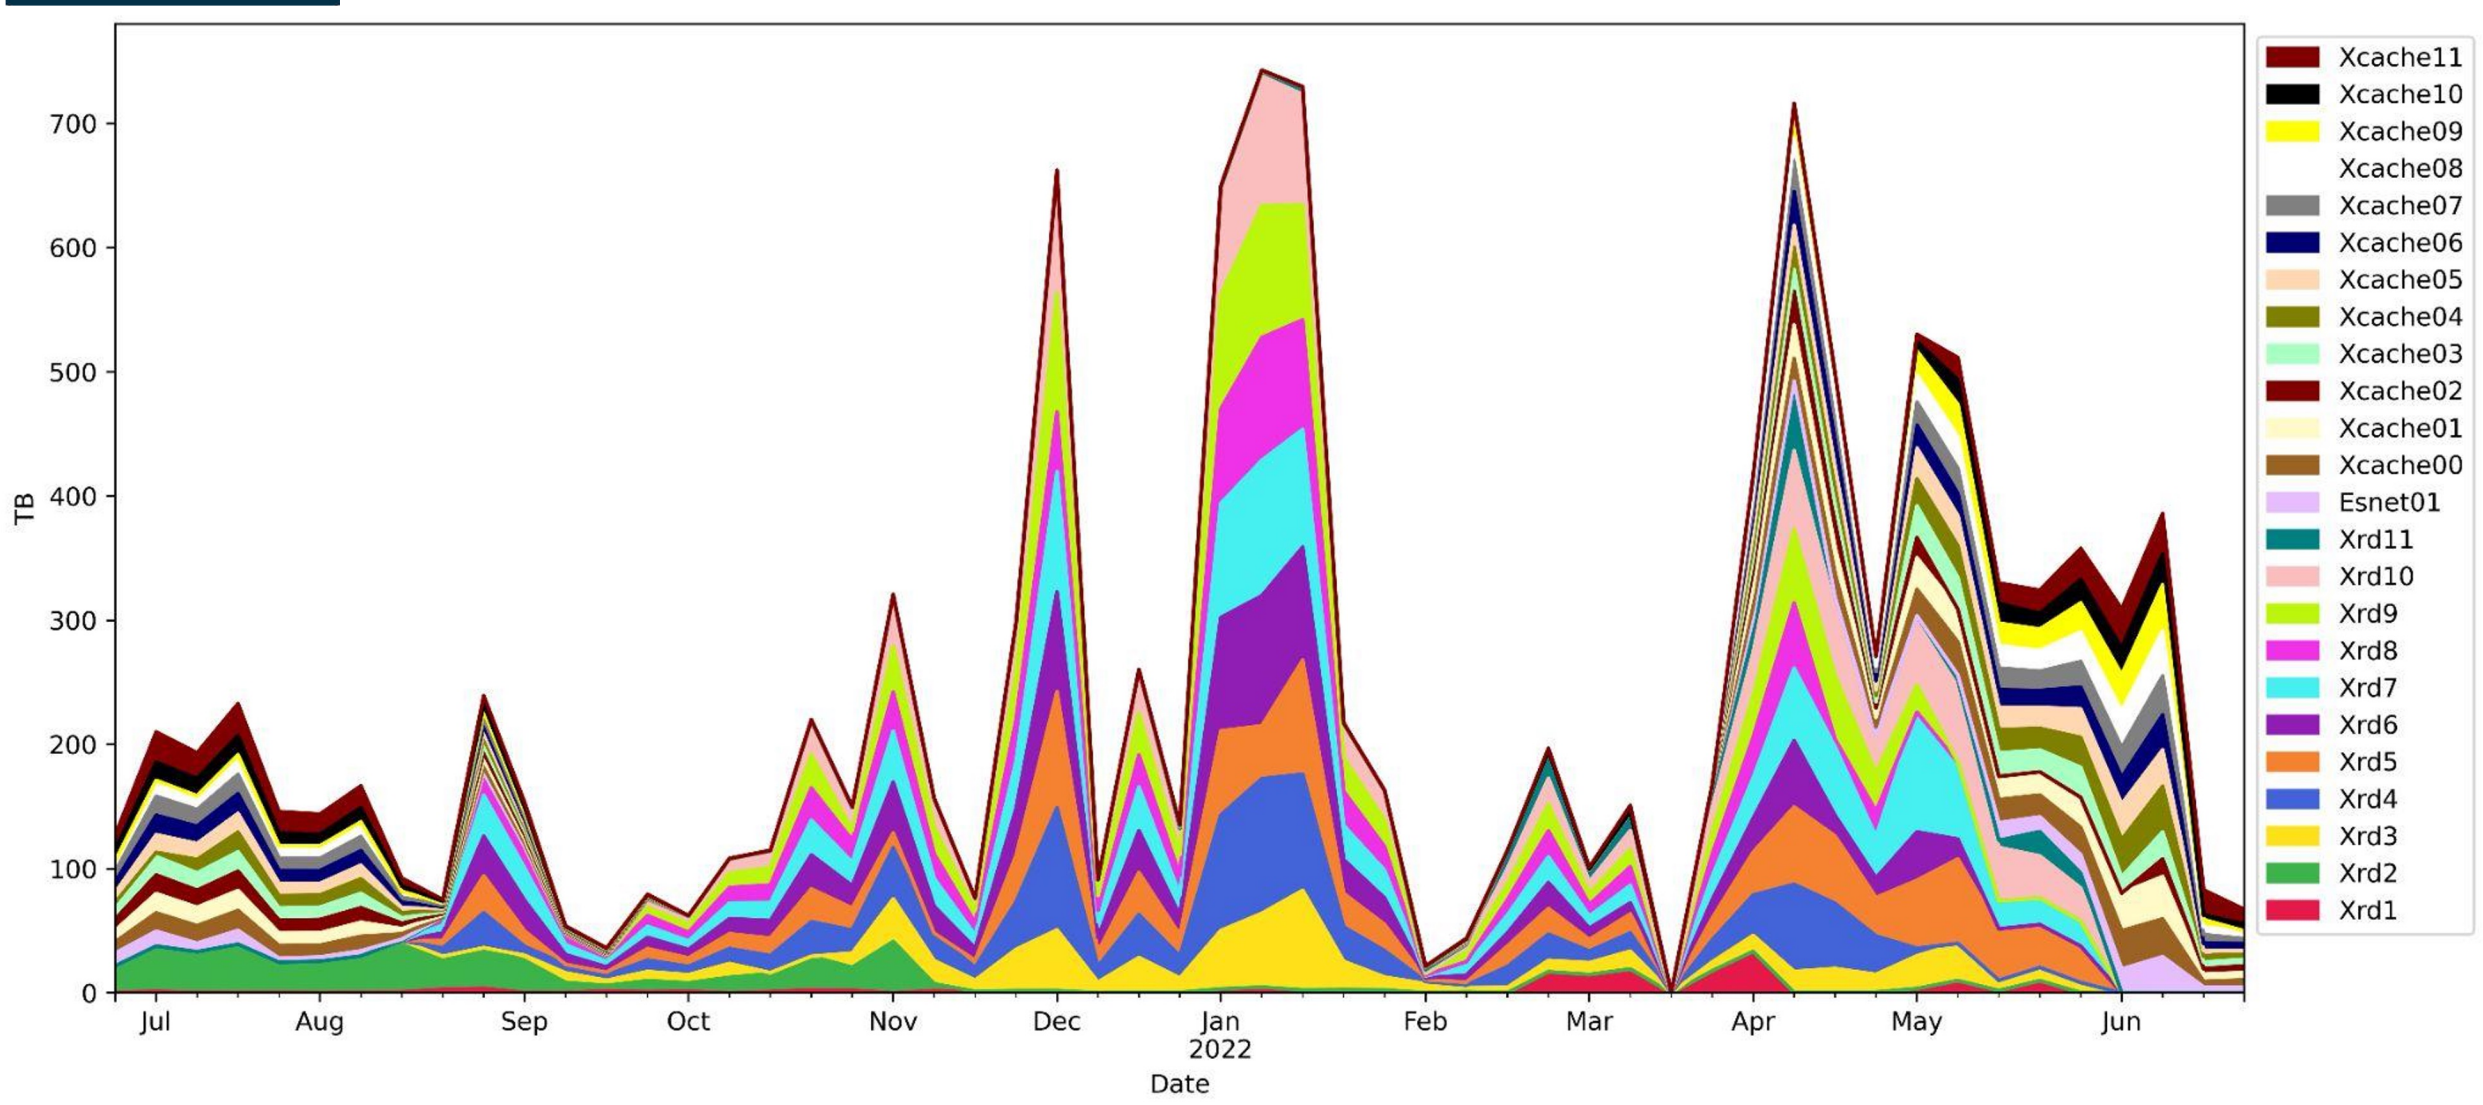

Daily Average Data Size per Access

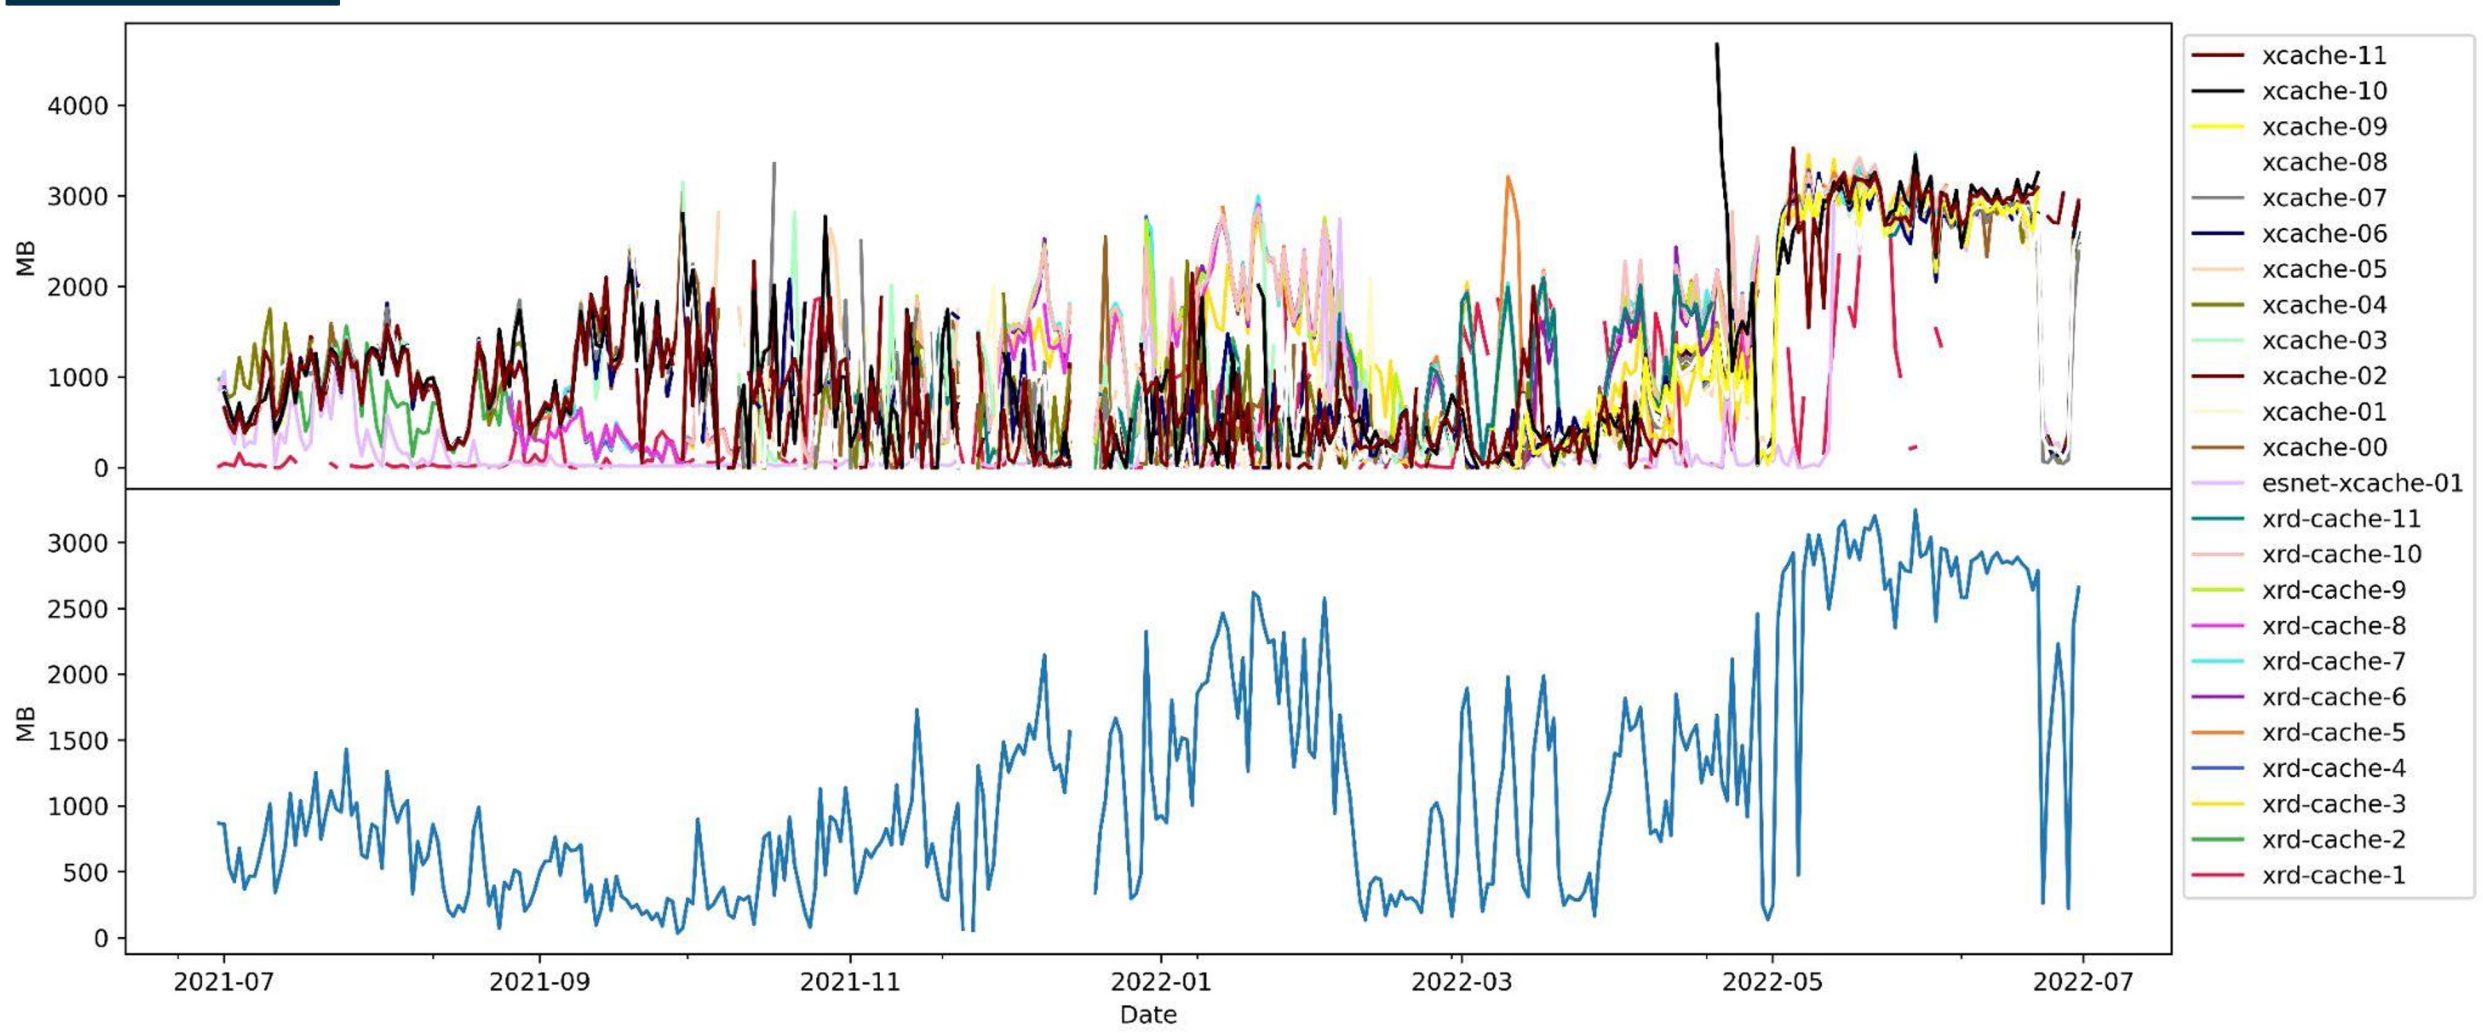

Weekly Access Data Size per Access

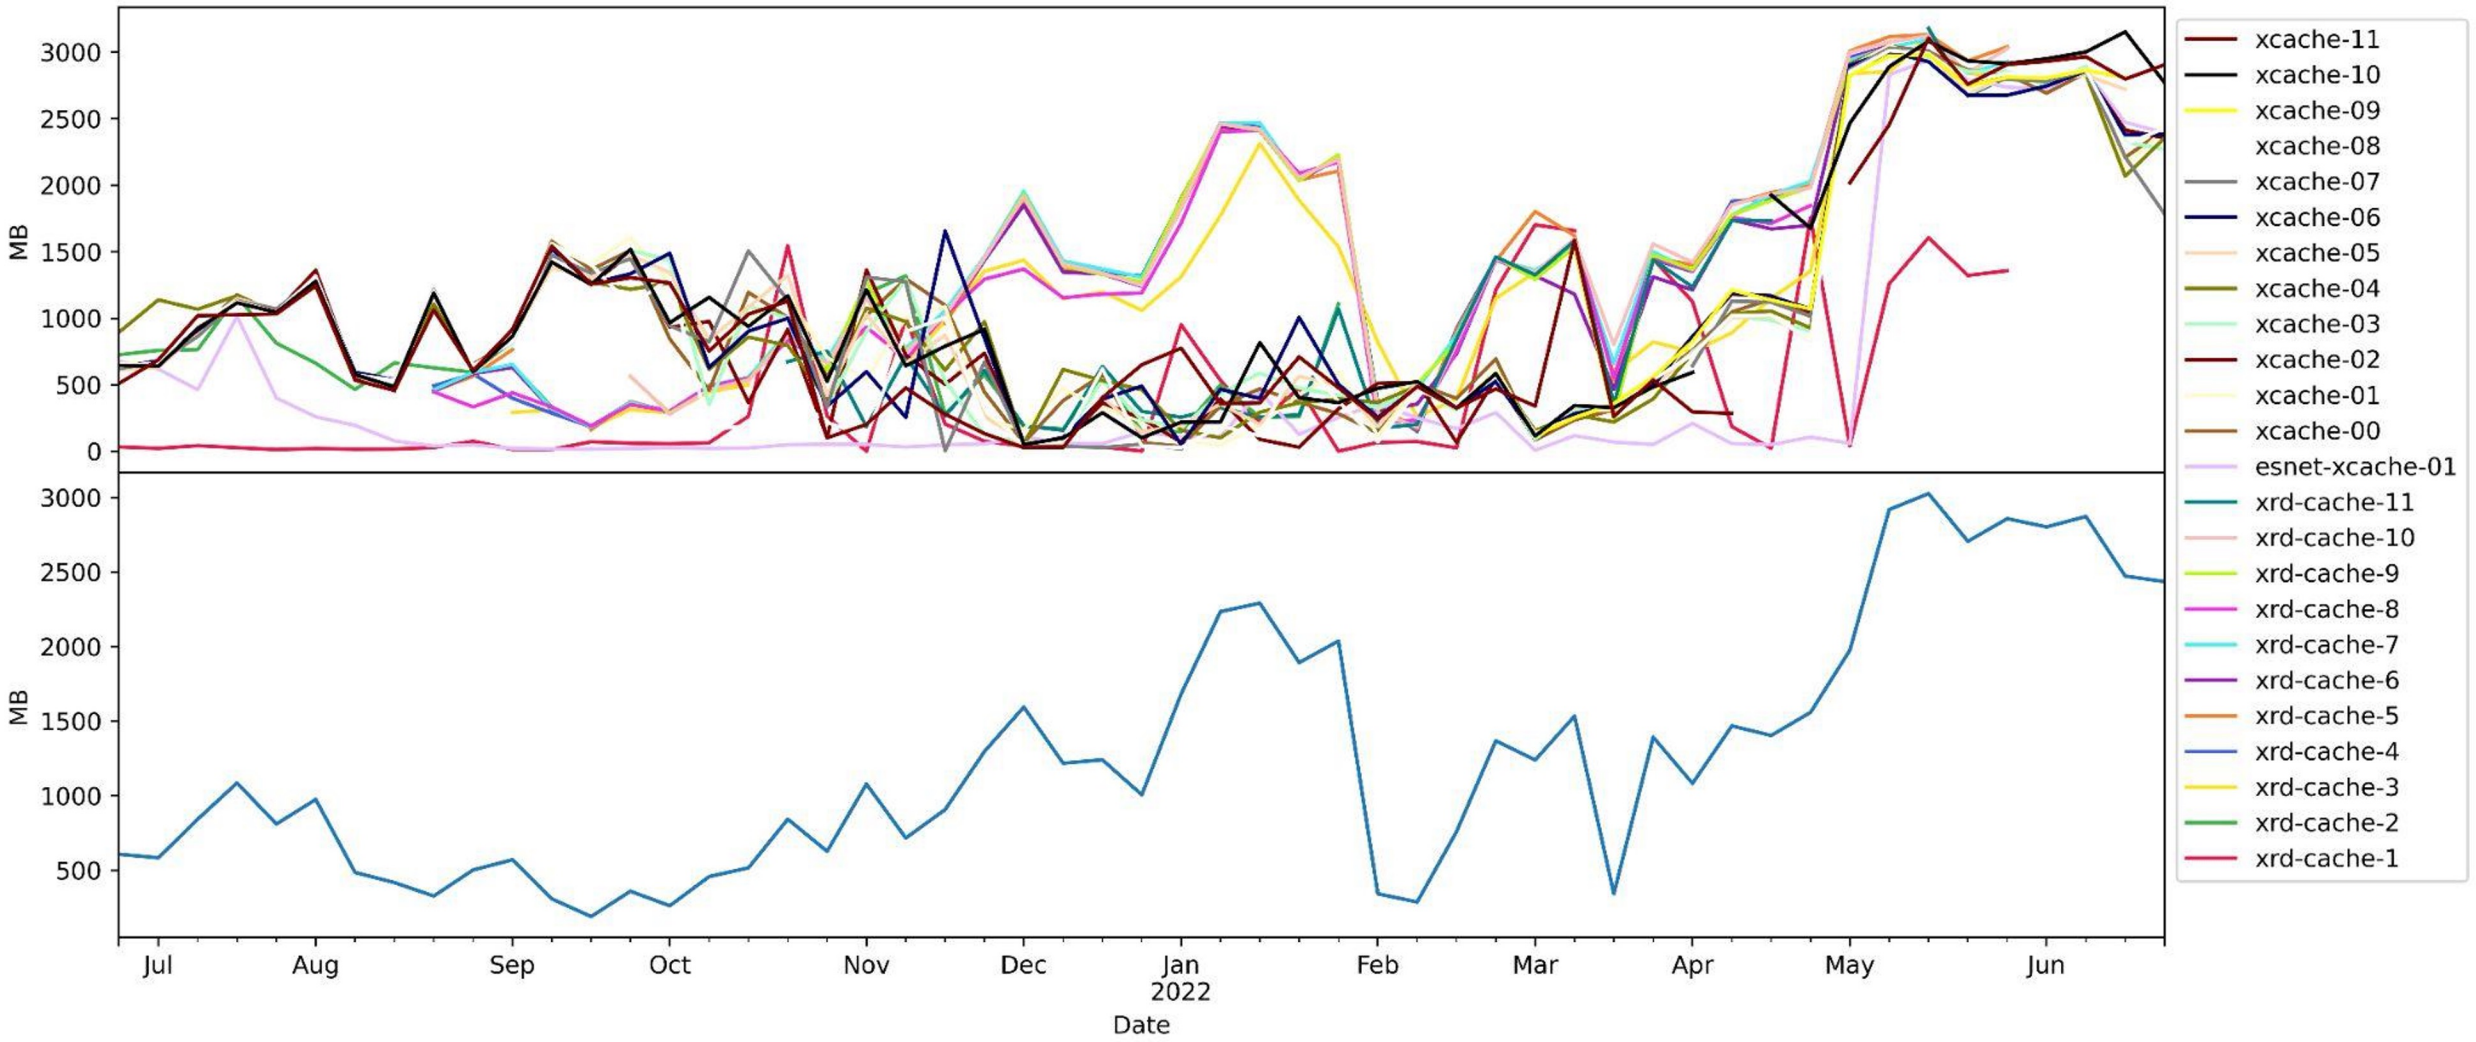

Daily number of cache hits

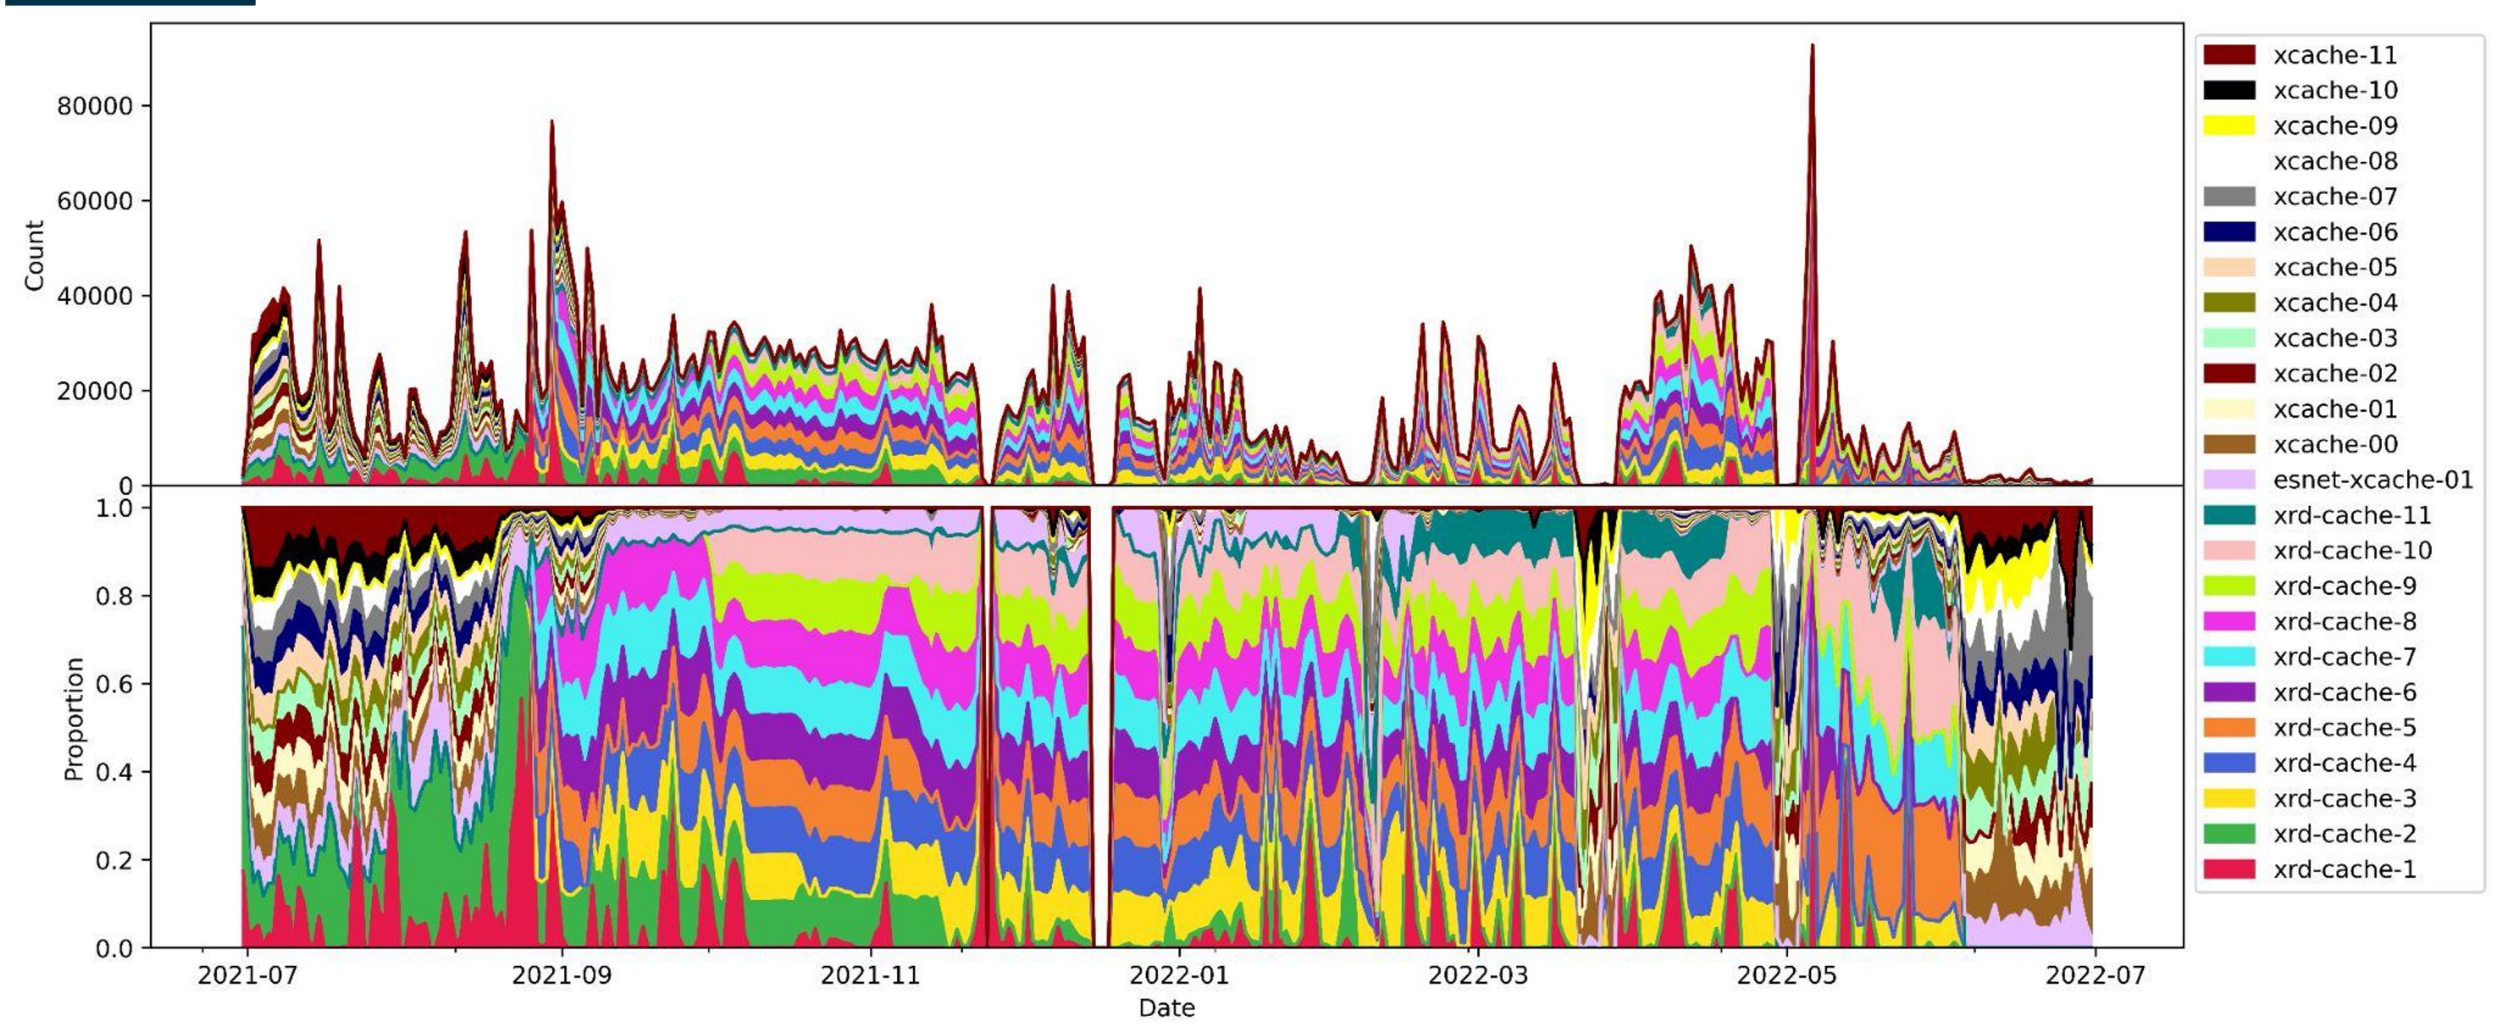

Weekly number of cache hits

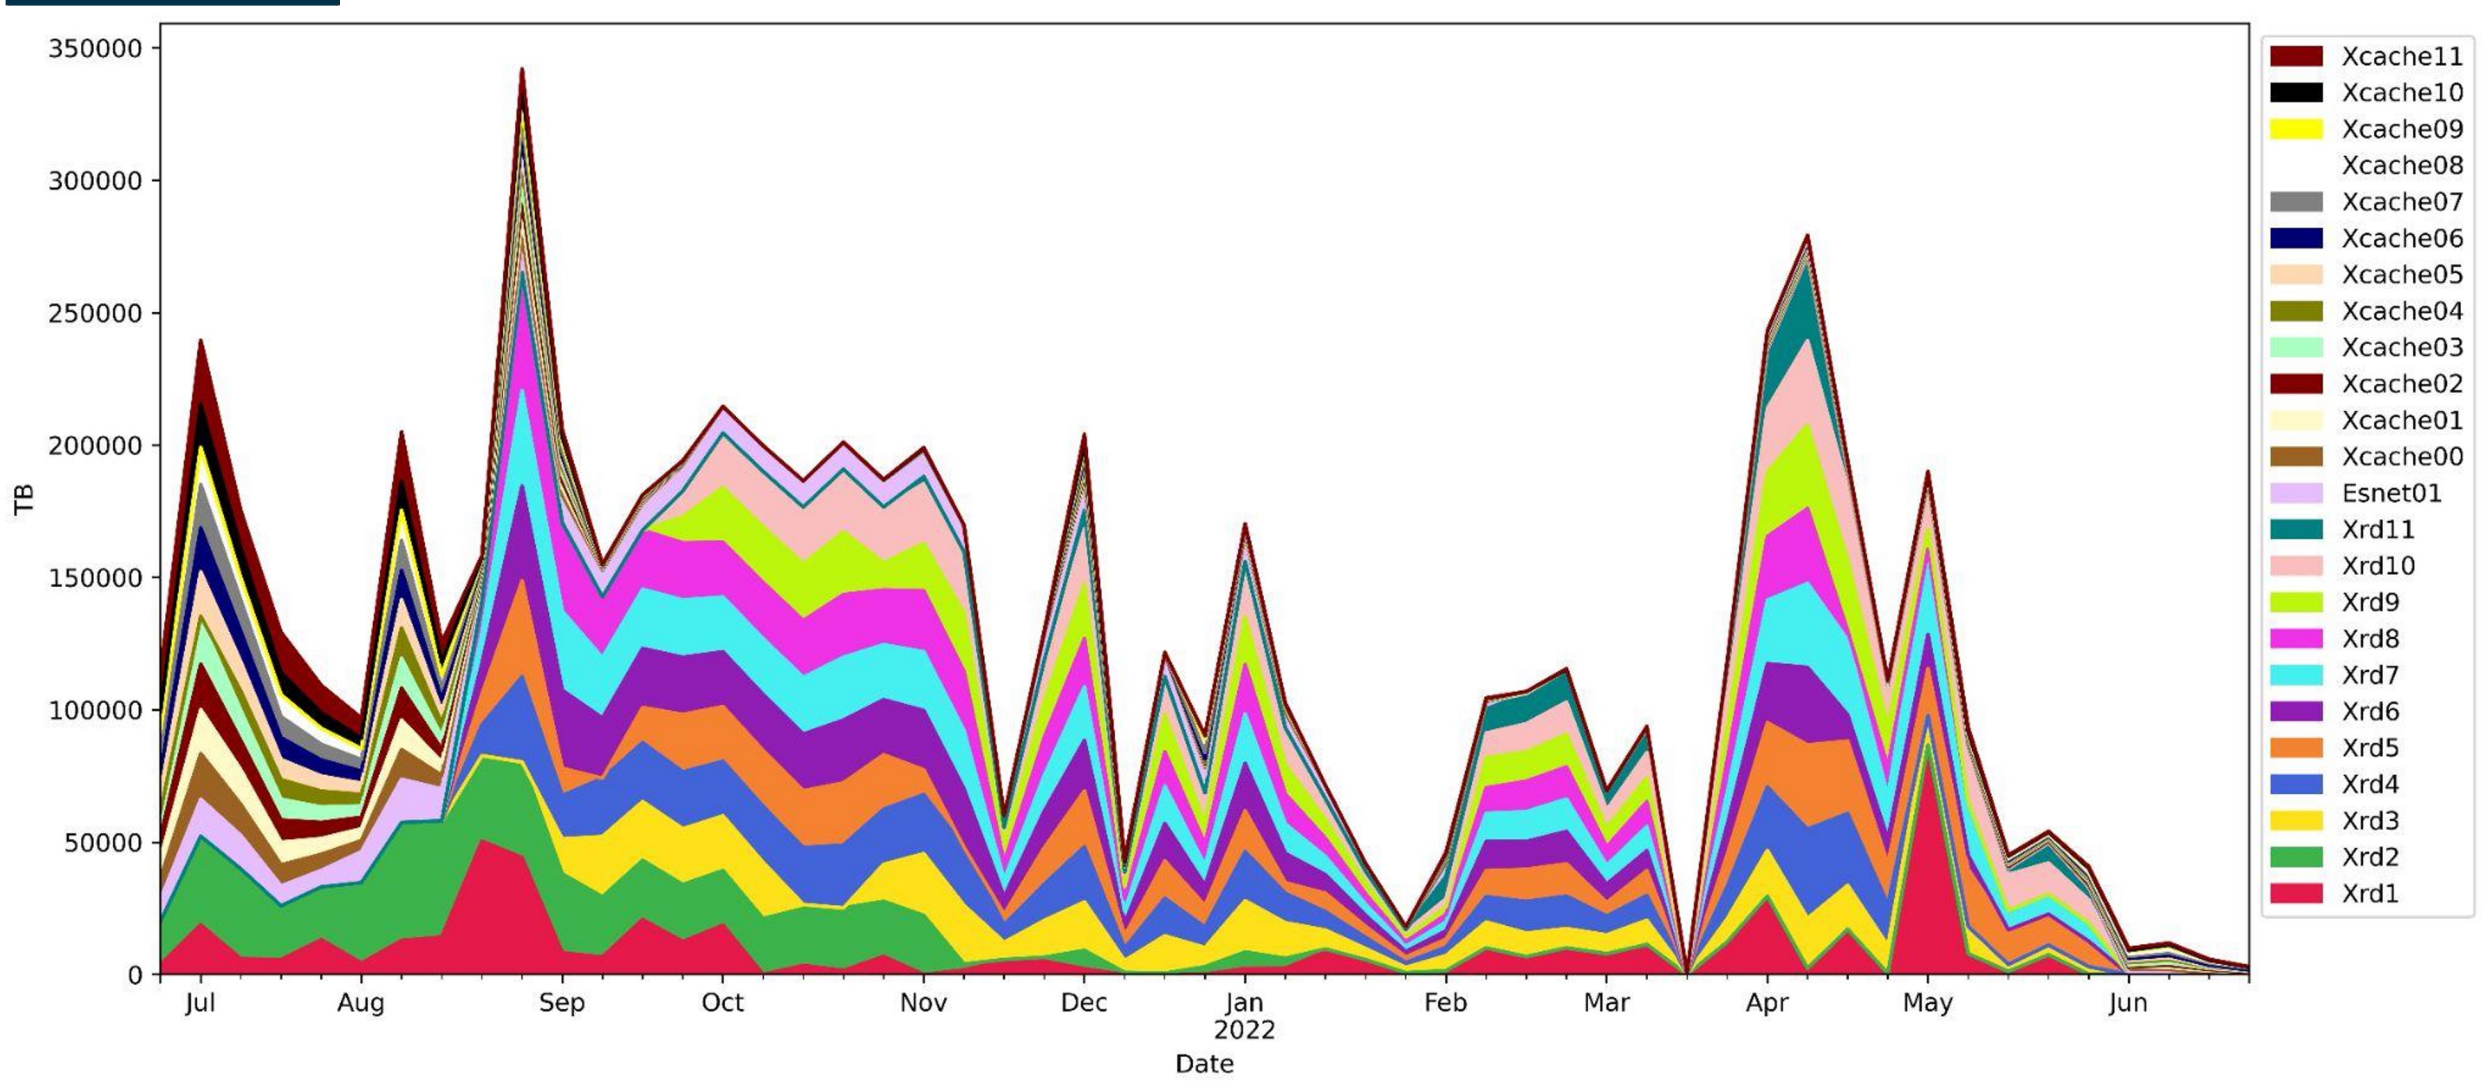

Daily number of cache misses

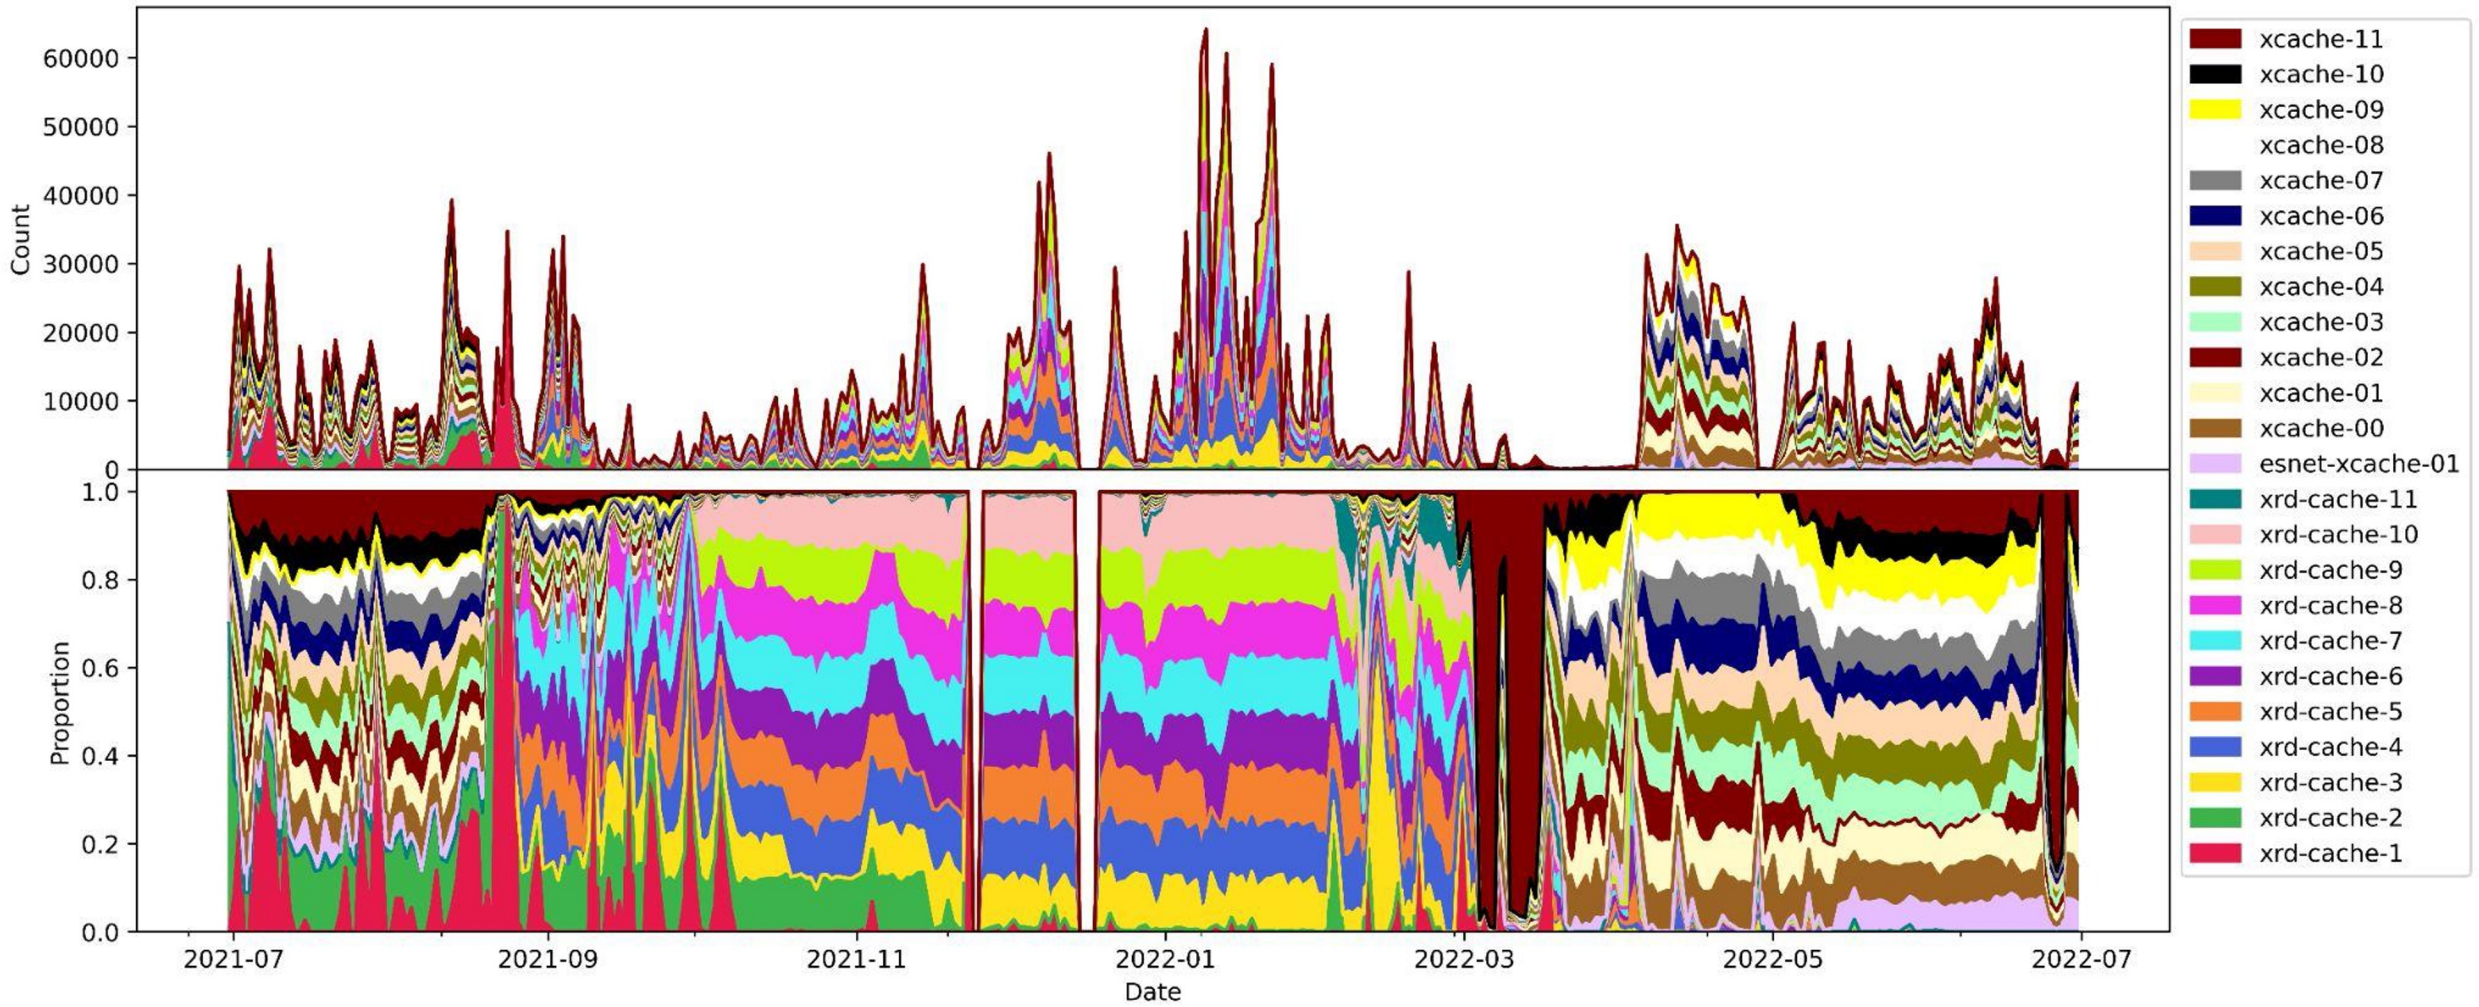

Weekly number of cache misses

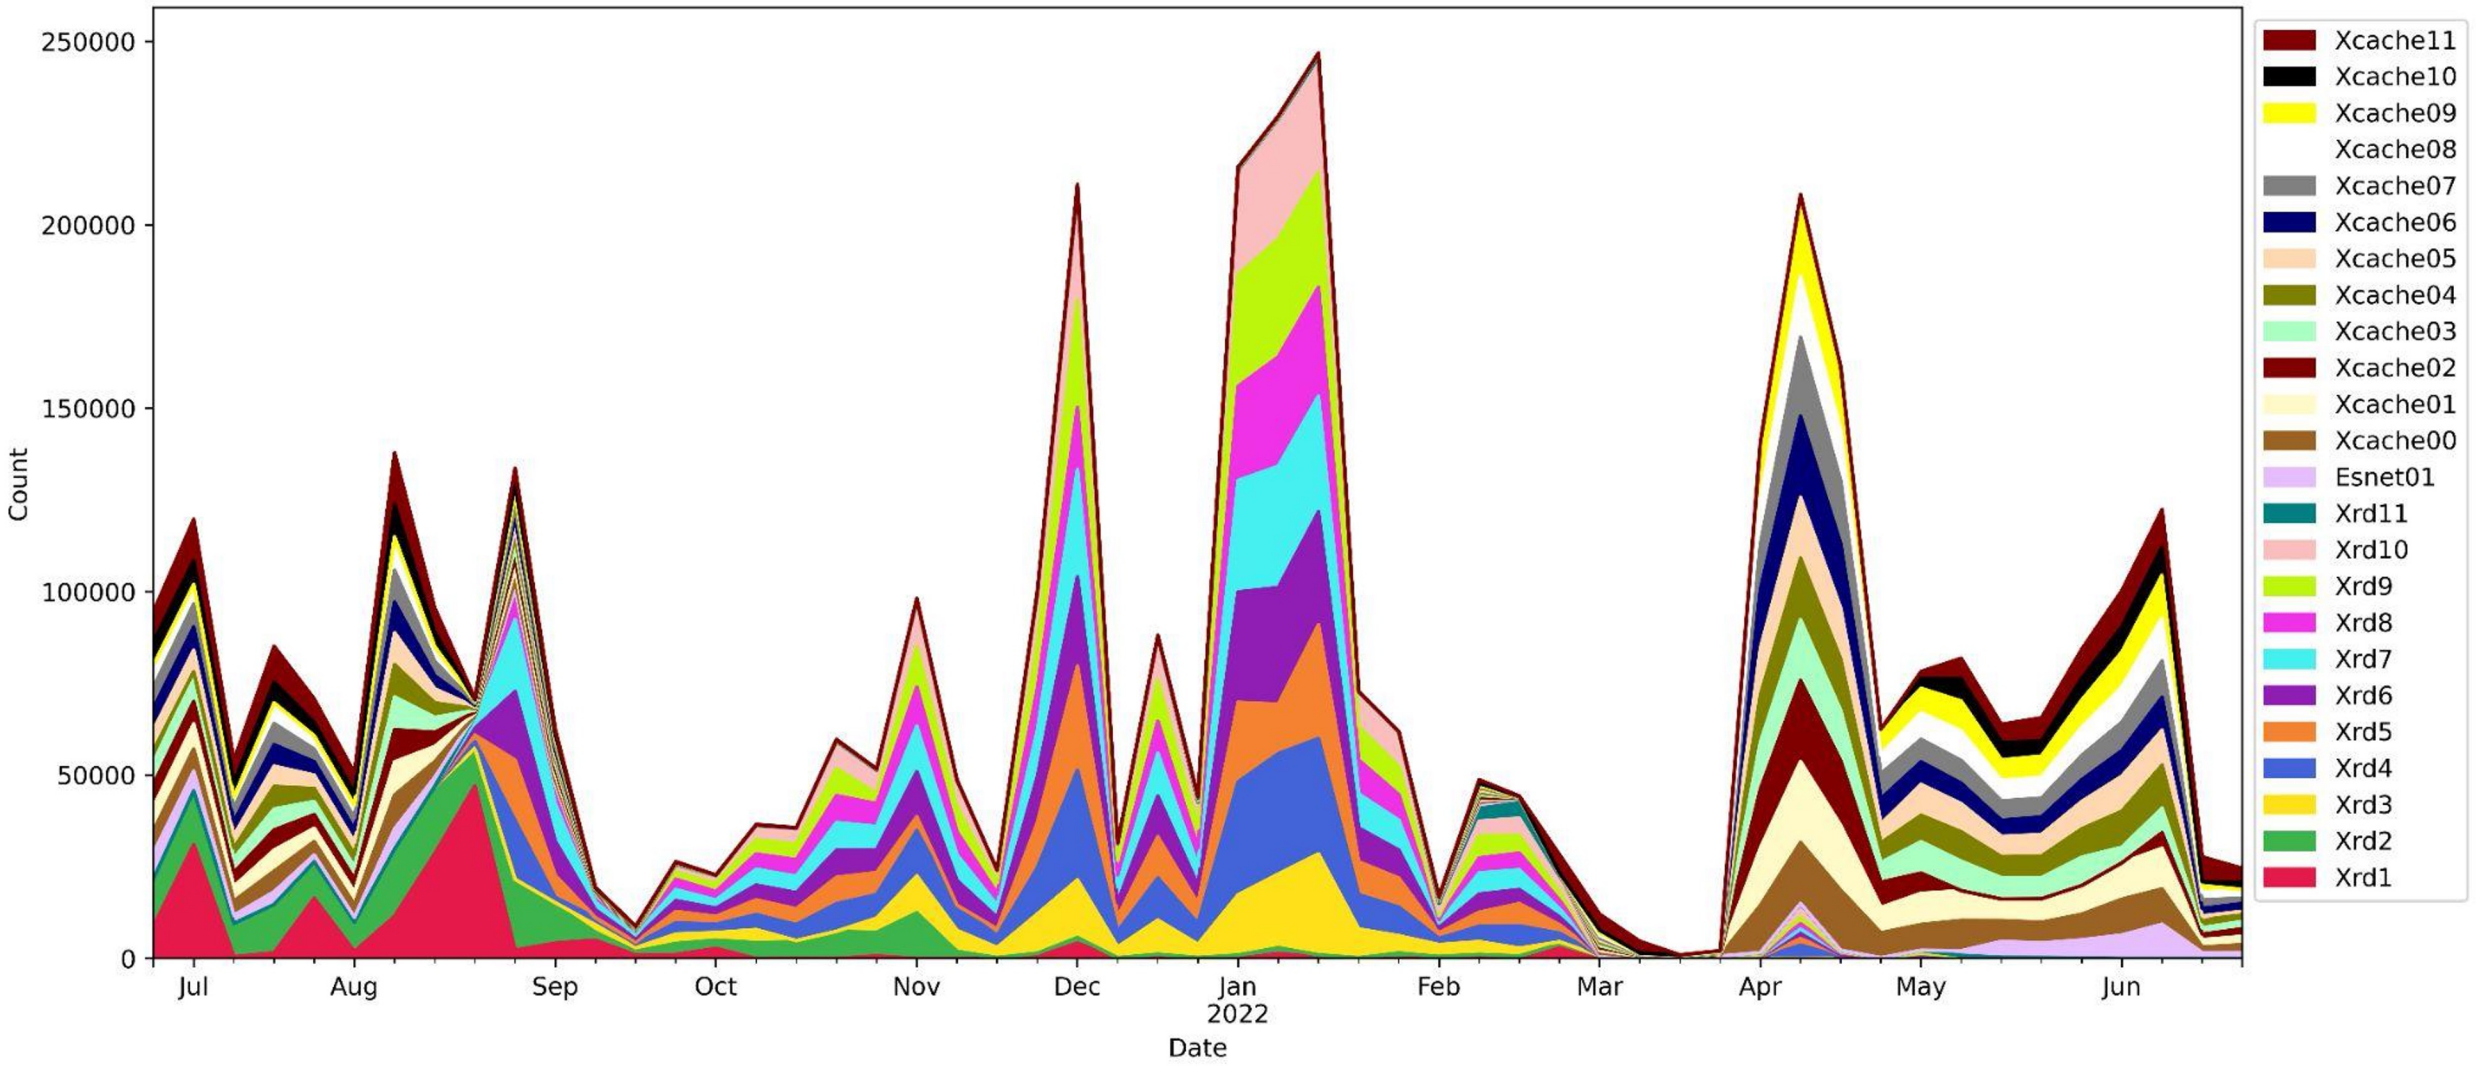

Daily proportion of number of cache hits and cache misses

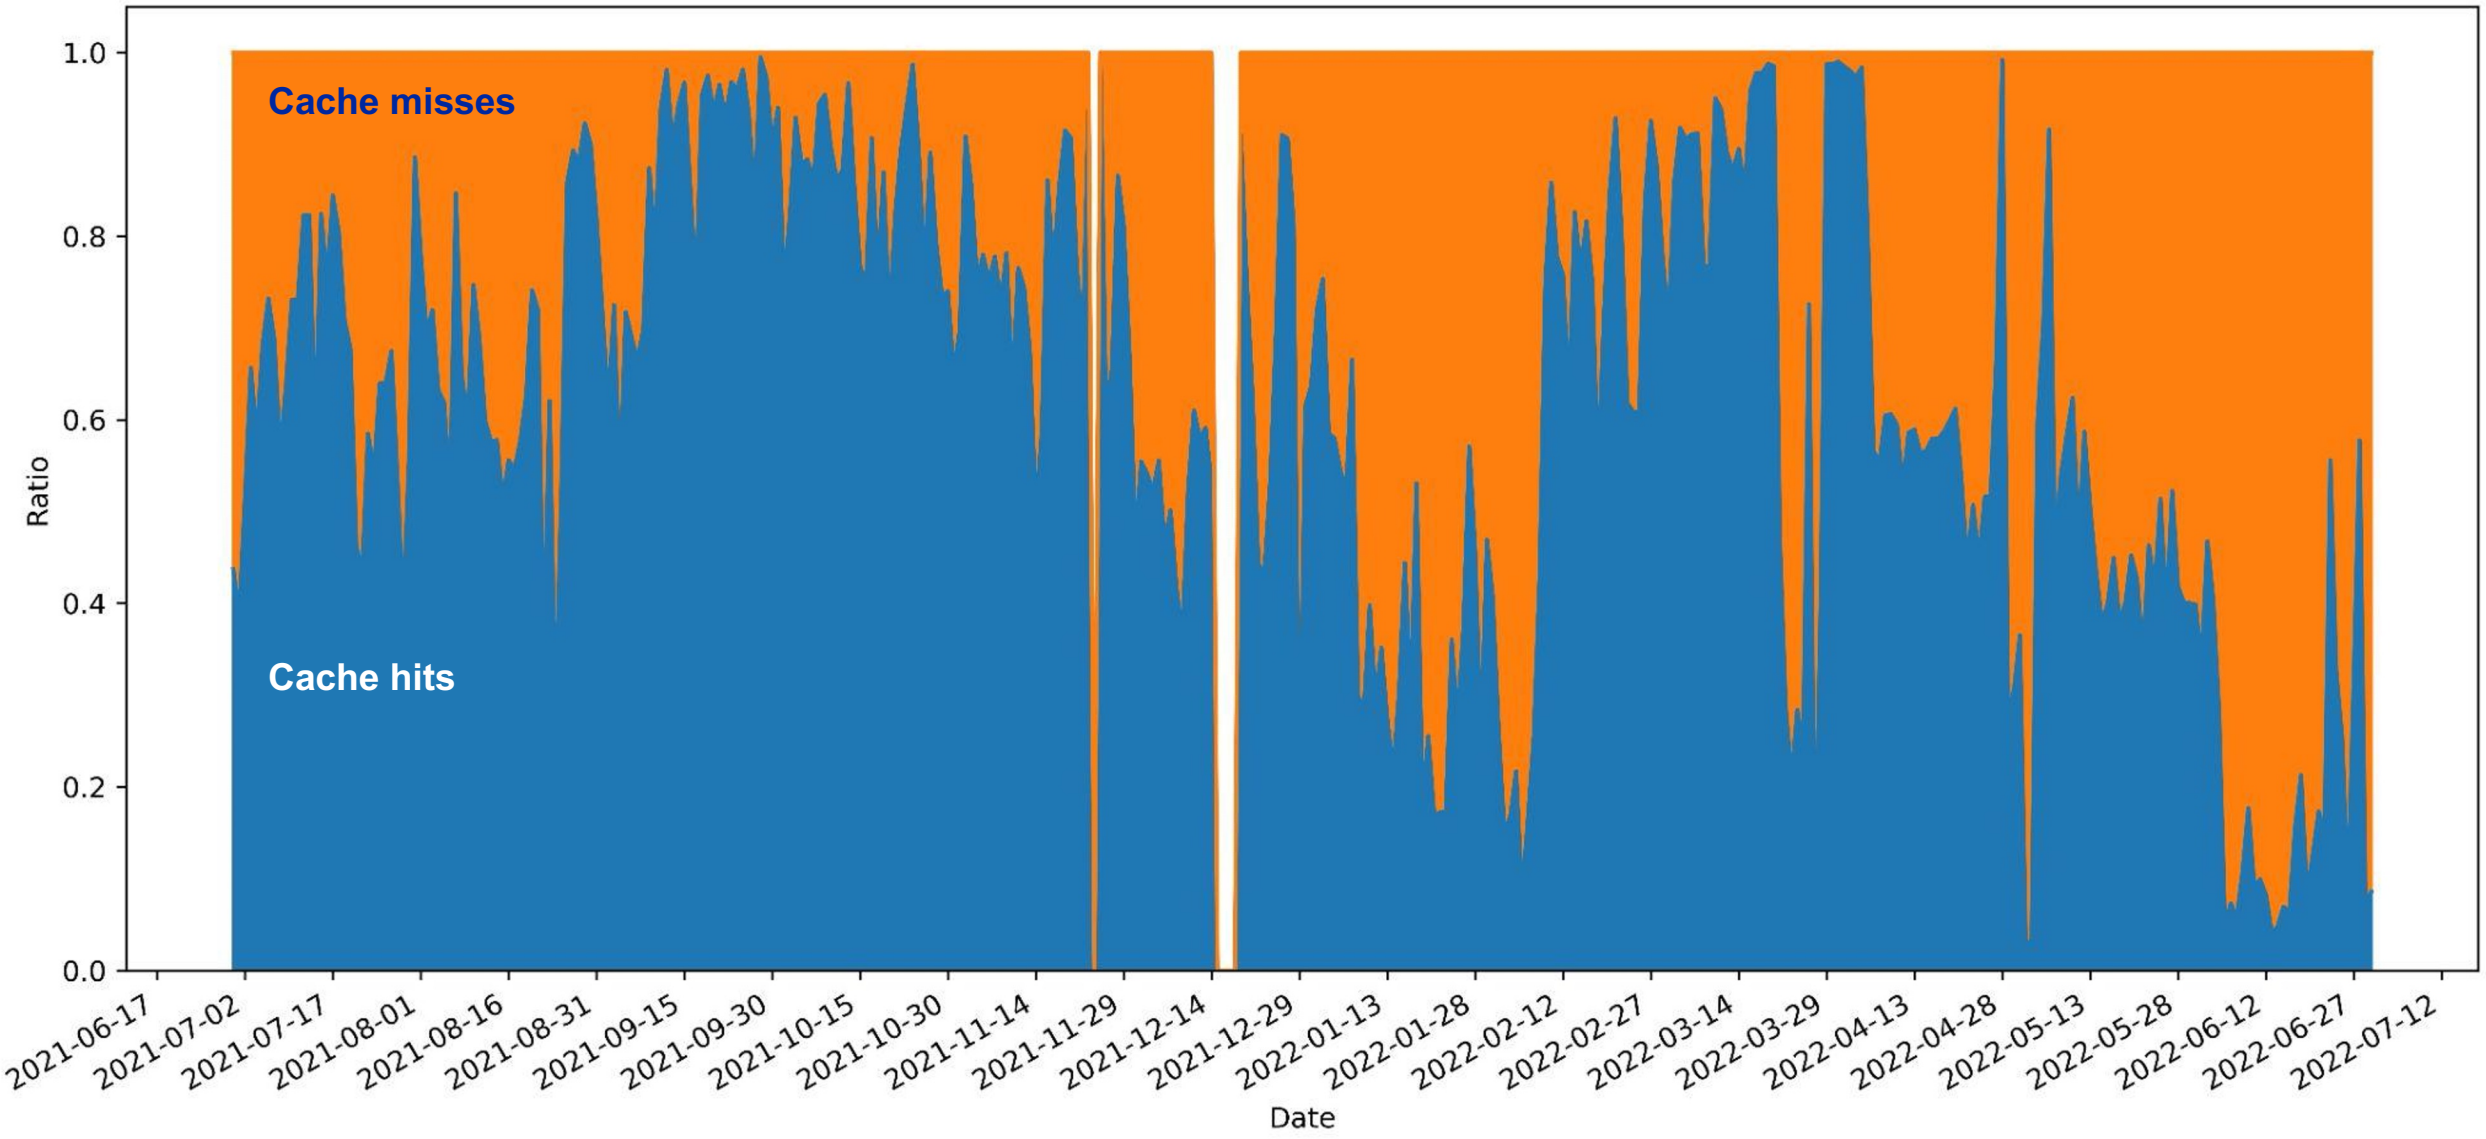

Weekly proportion of number of cache hits

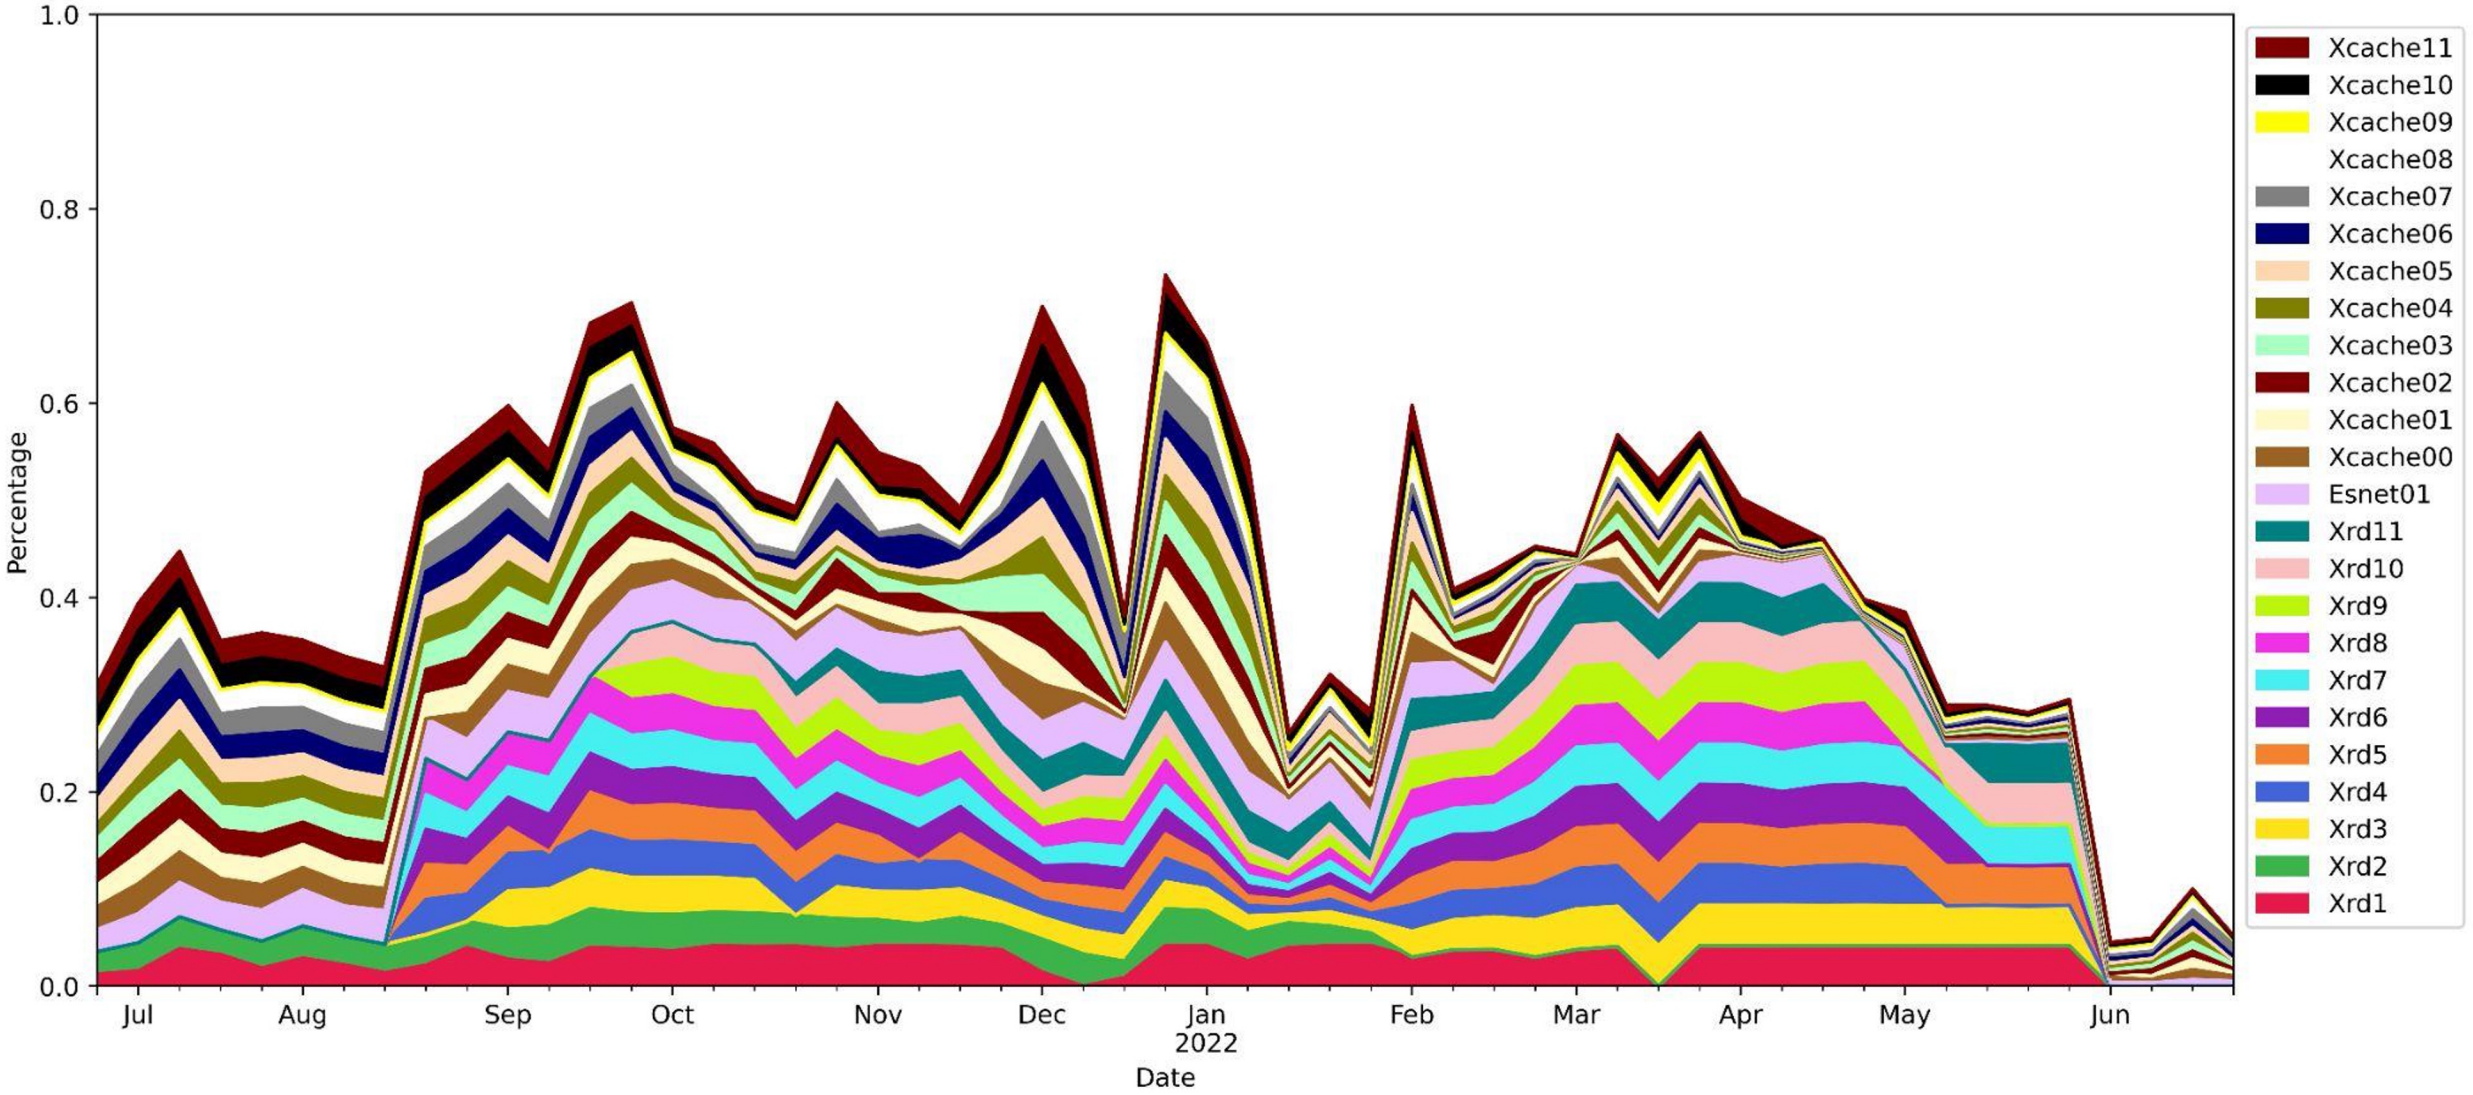

Daily network demand reduction rate (traffic frequency)

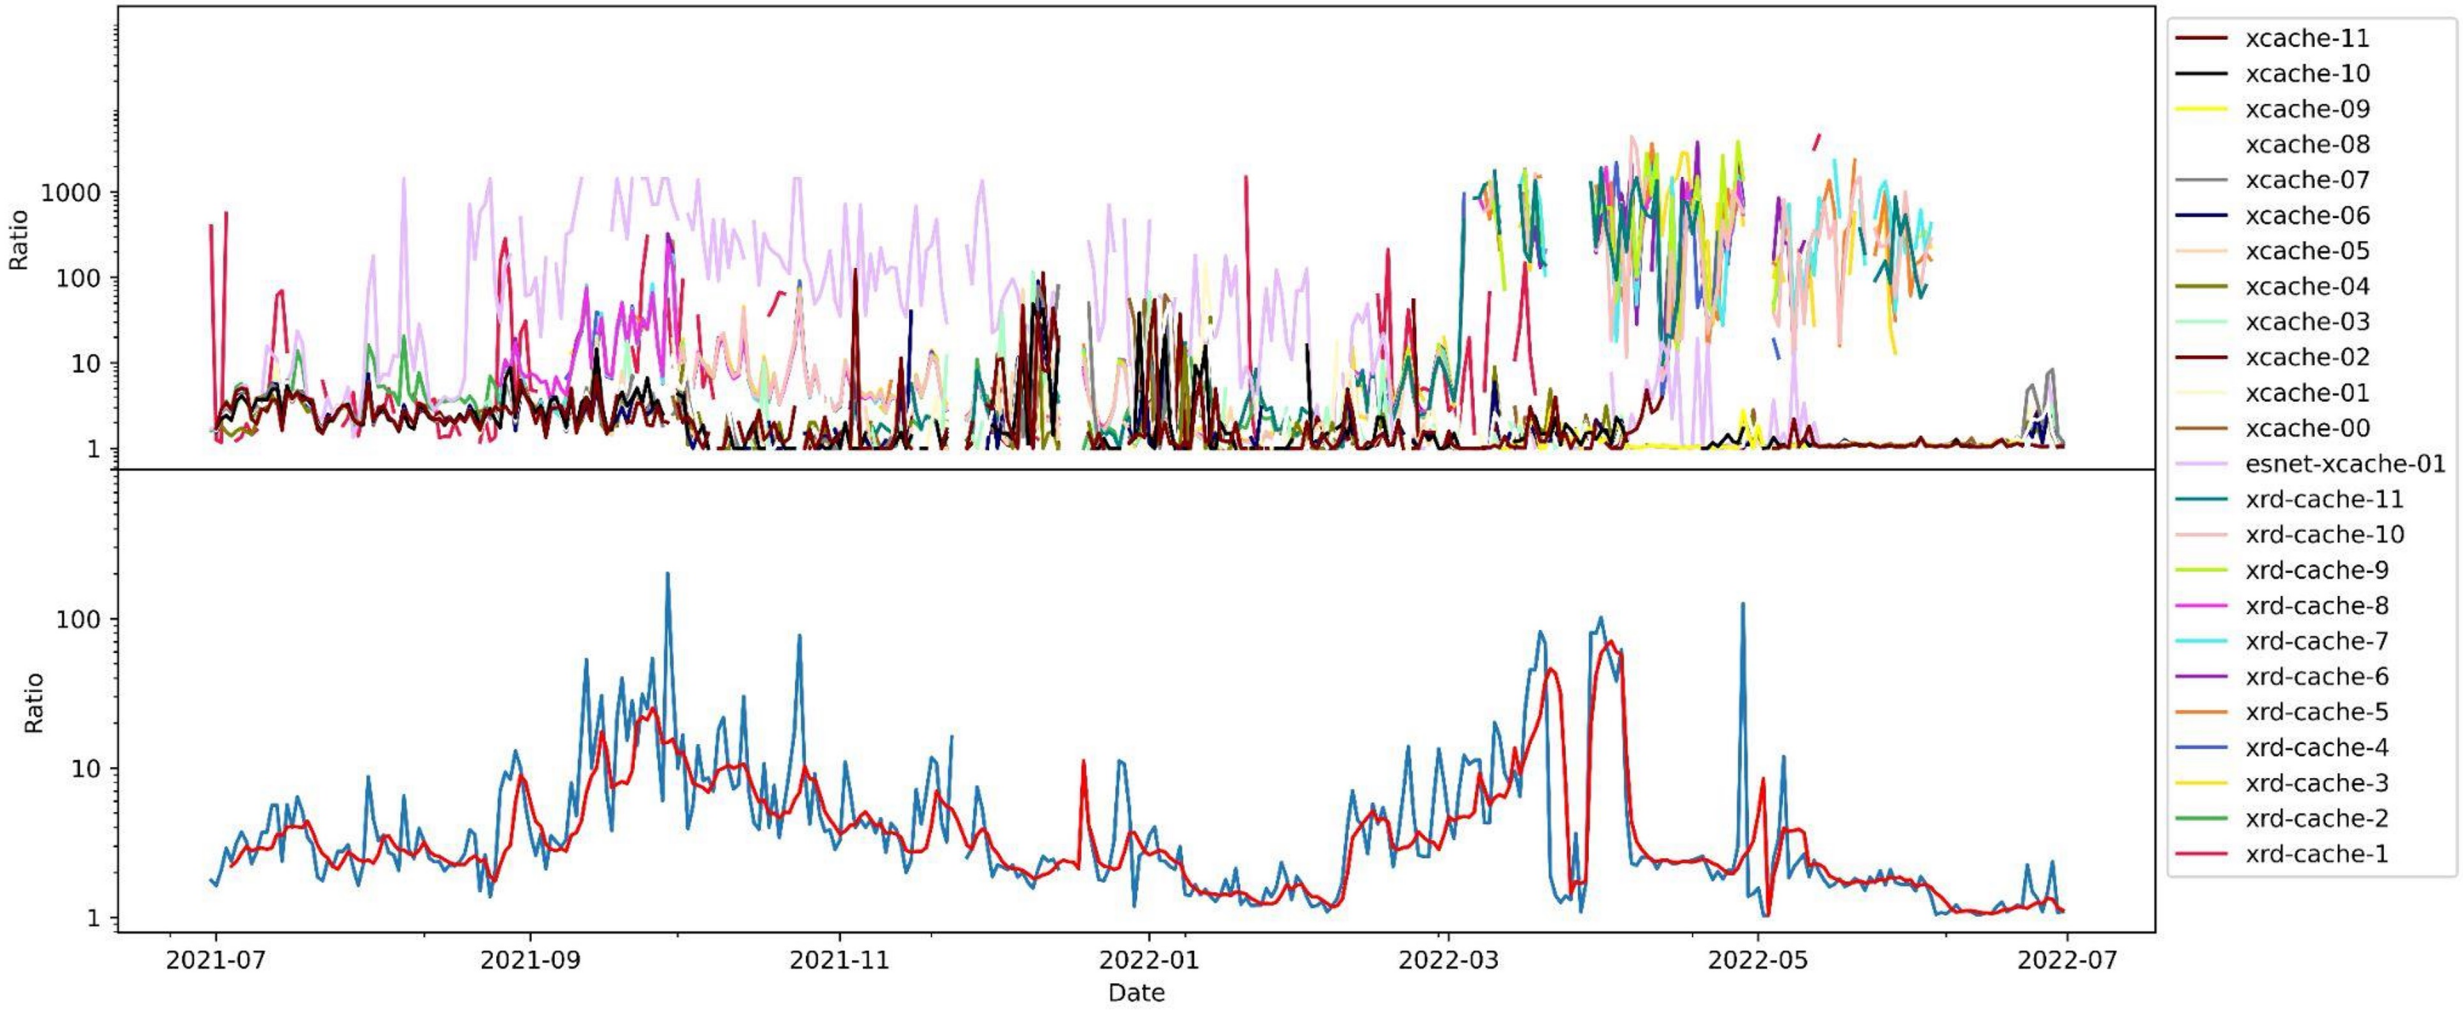

Avg traffic frequency reduction: 2.66

Daily cache hits volume (shared data size)

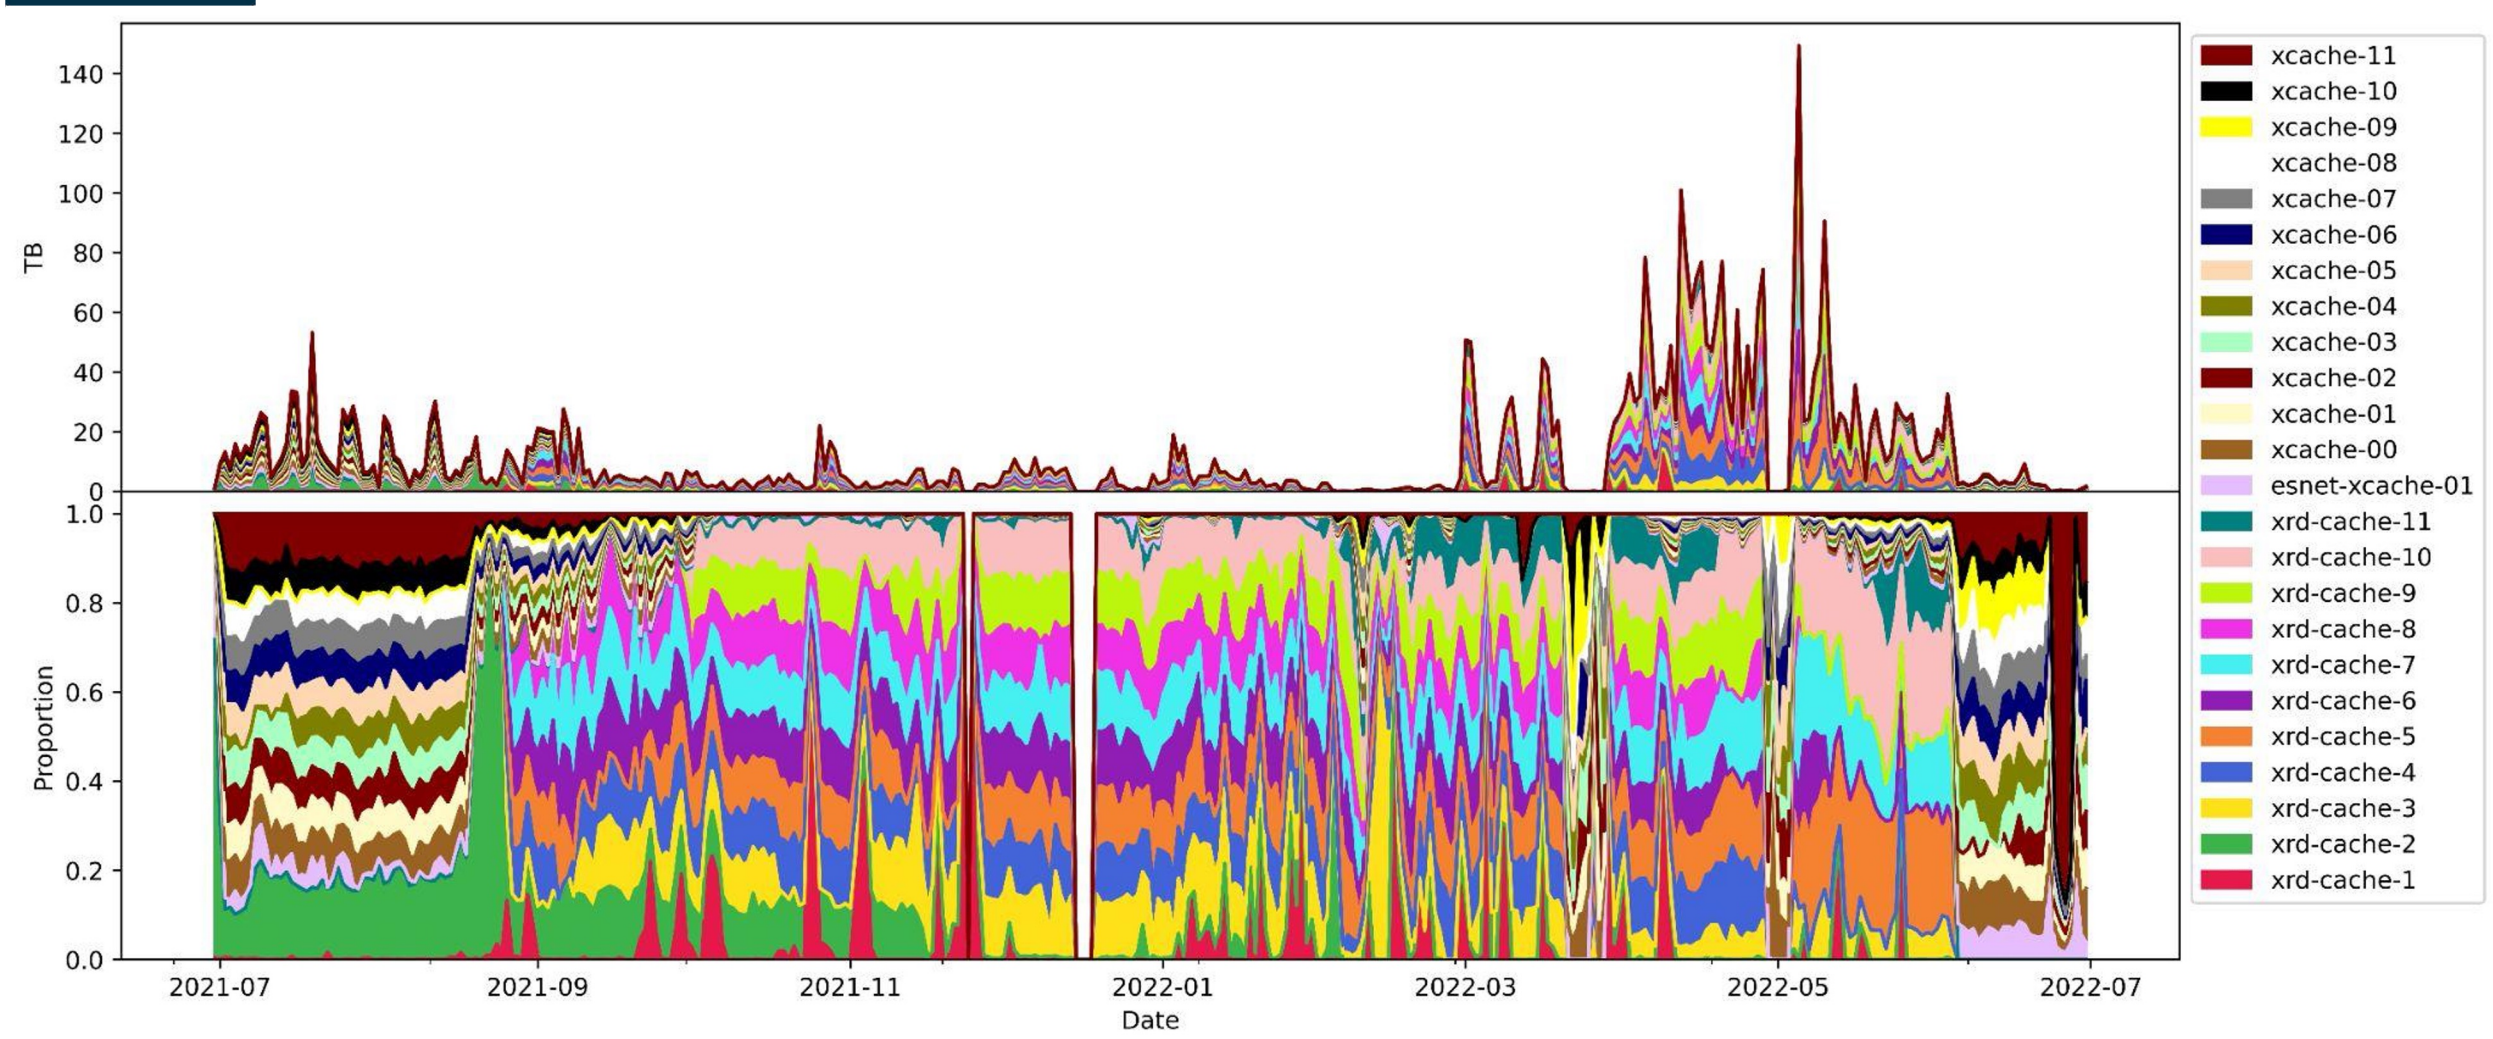

Weekly cache hits volume (shared data size)

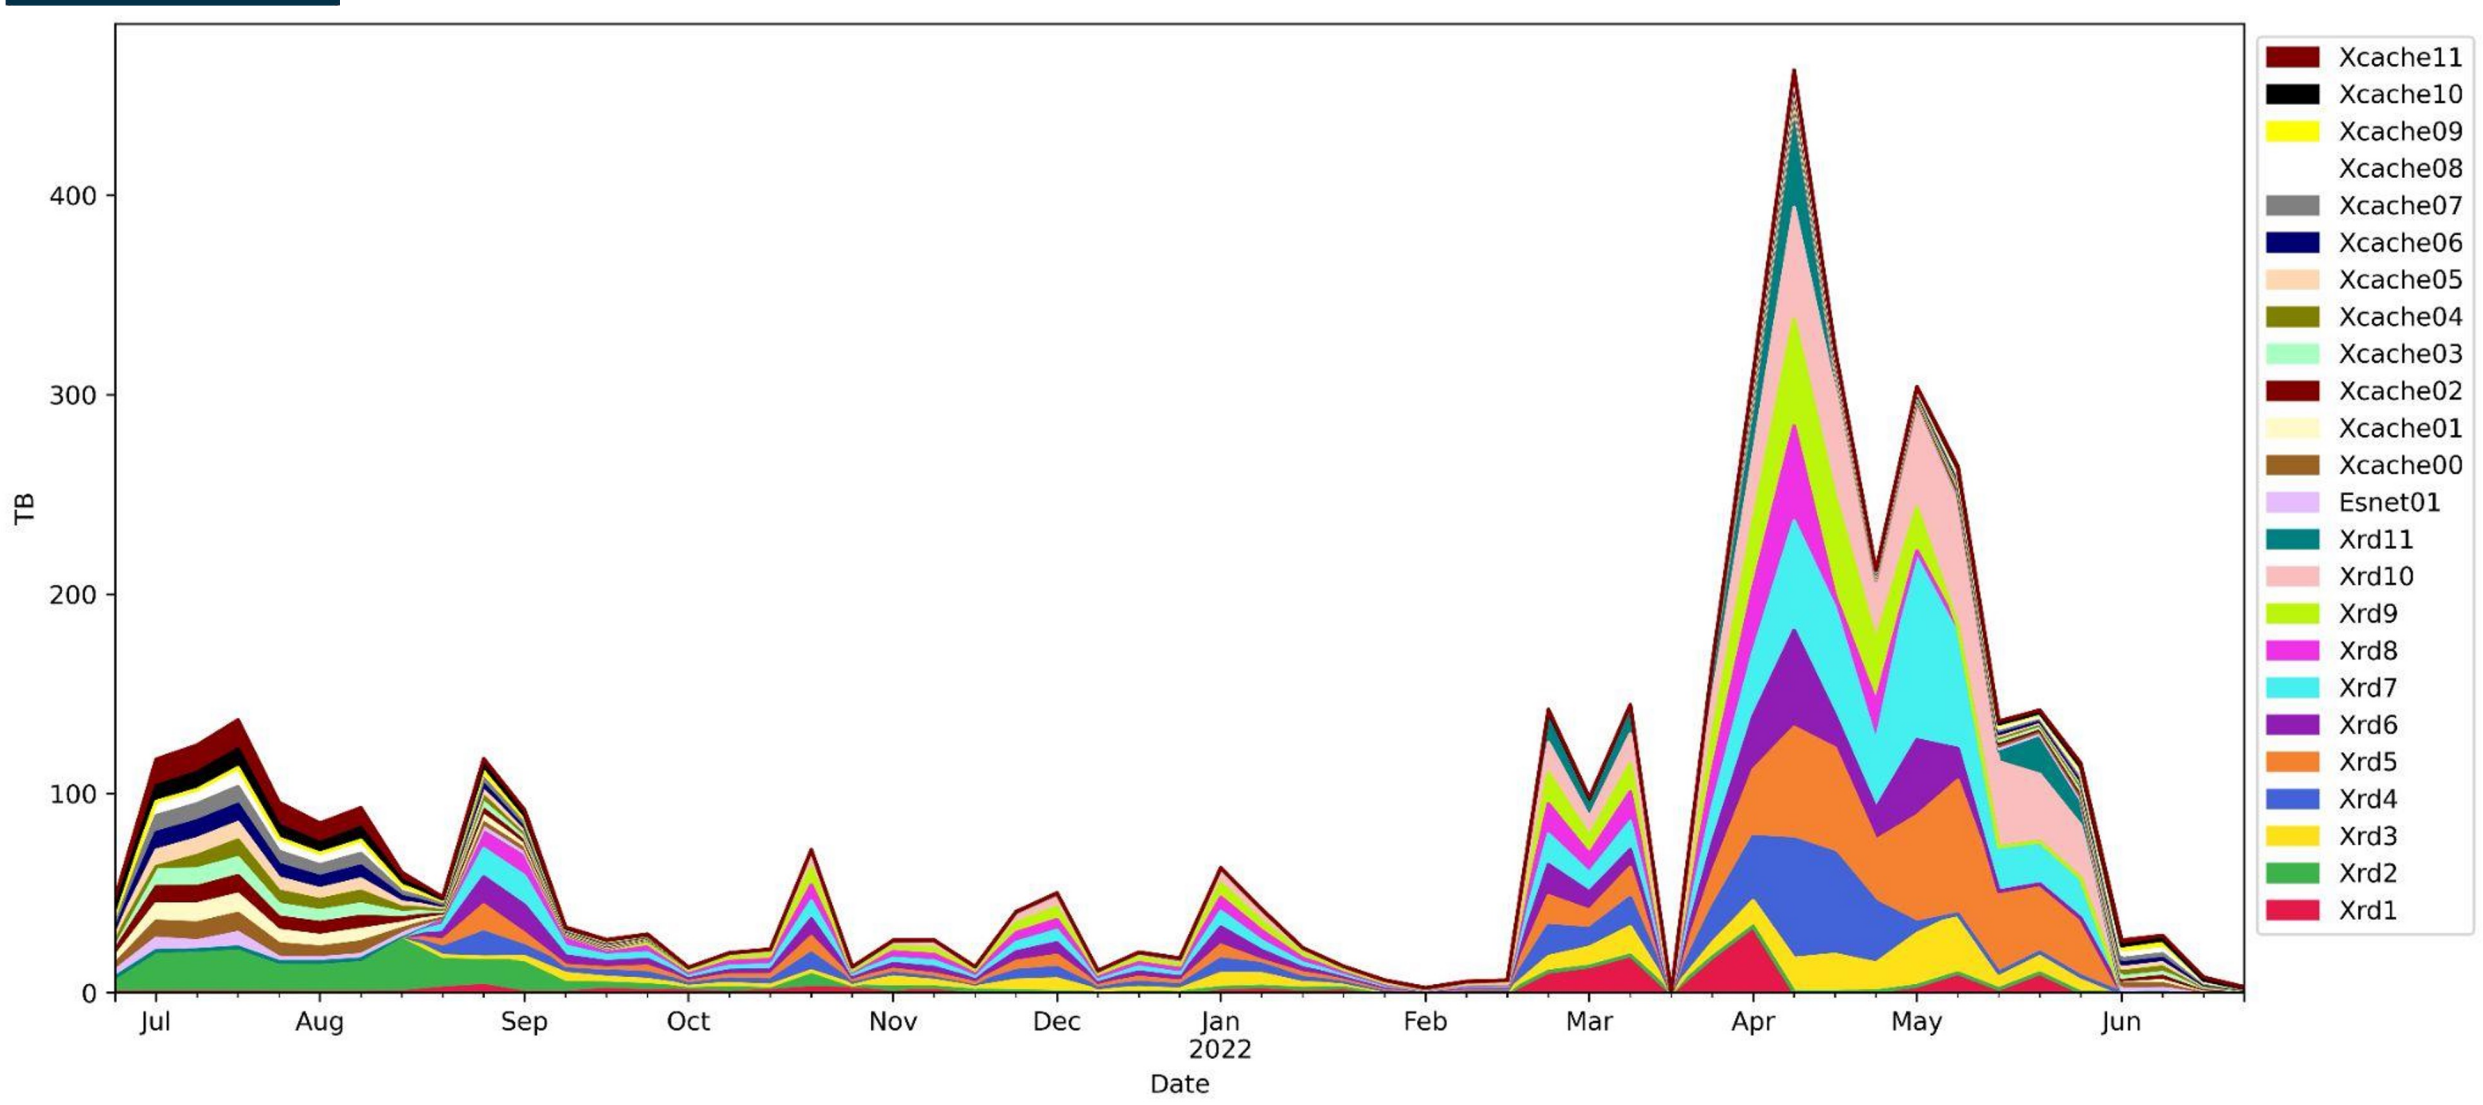

Daily cache misses volume (data transfer size)

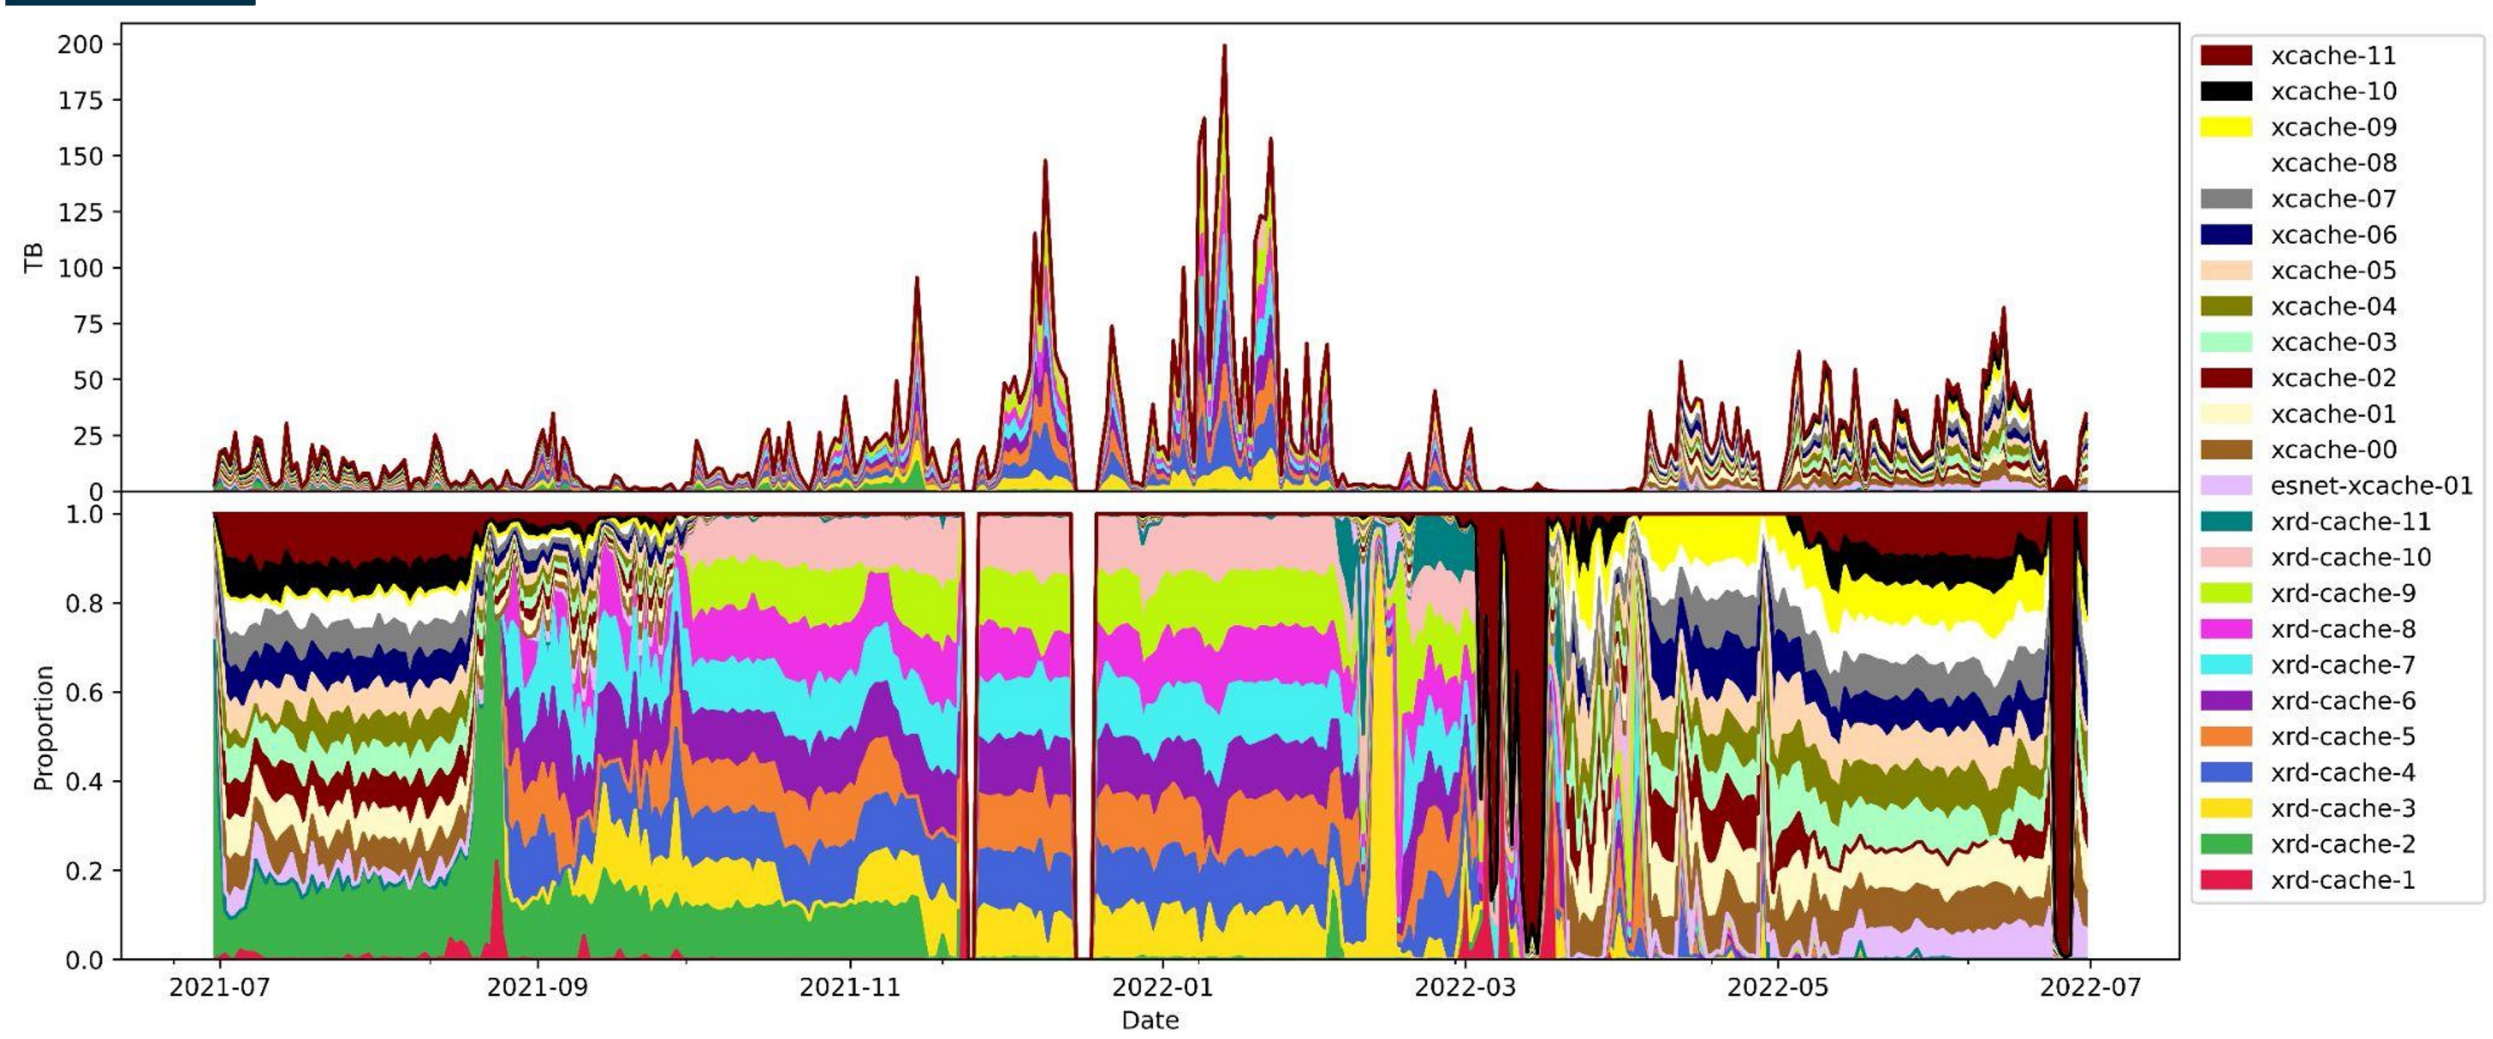

Weekly cache misses volume (data transfer size)

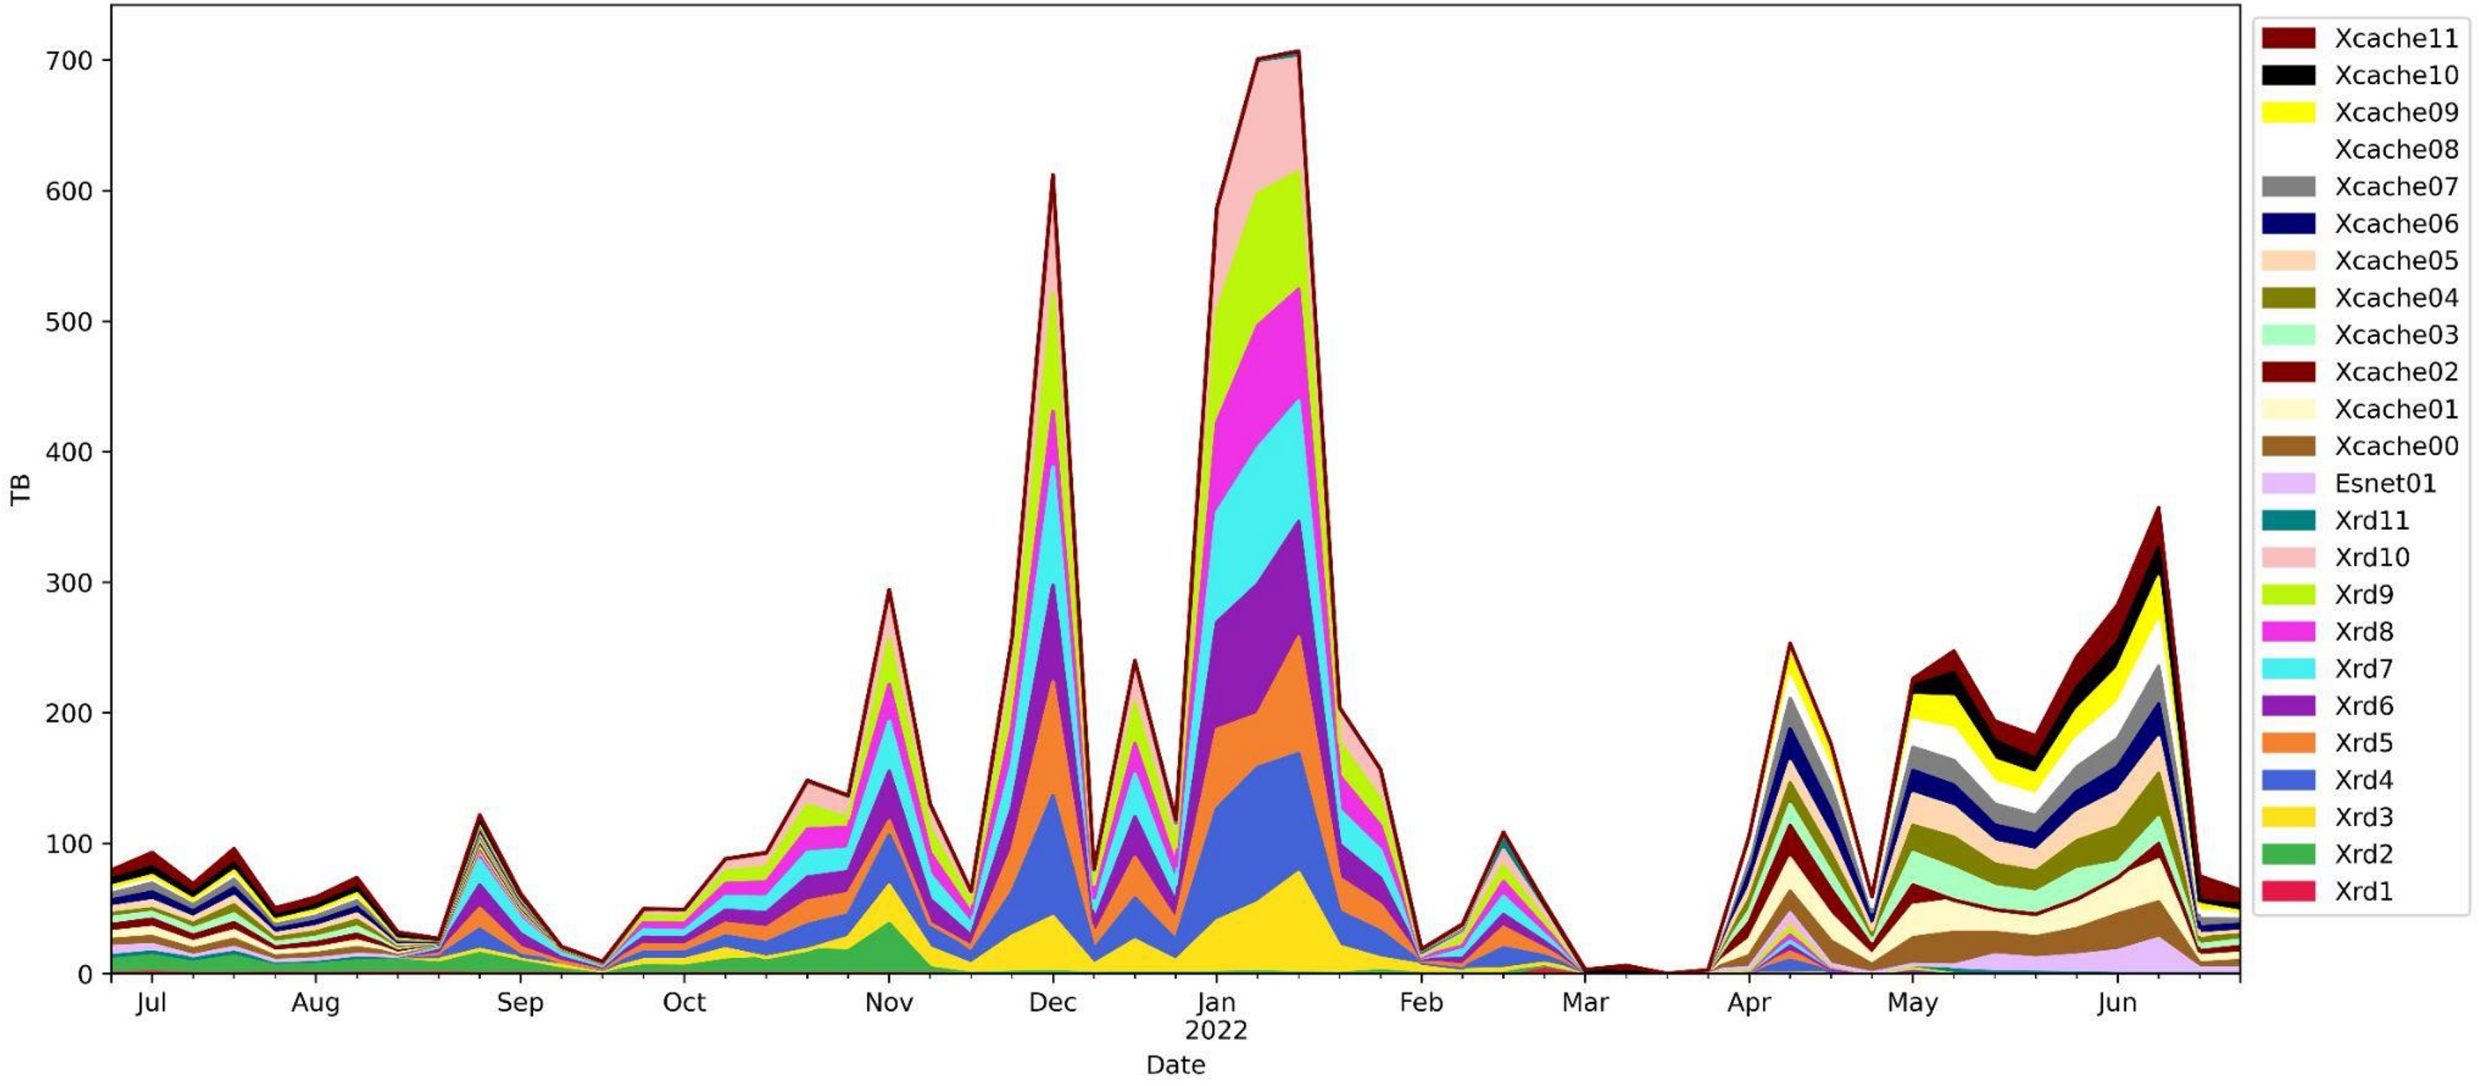

Daily proportion of cache hits volume and cache misses volume

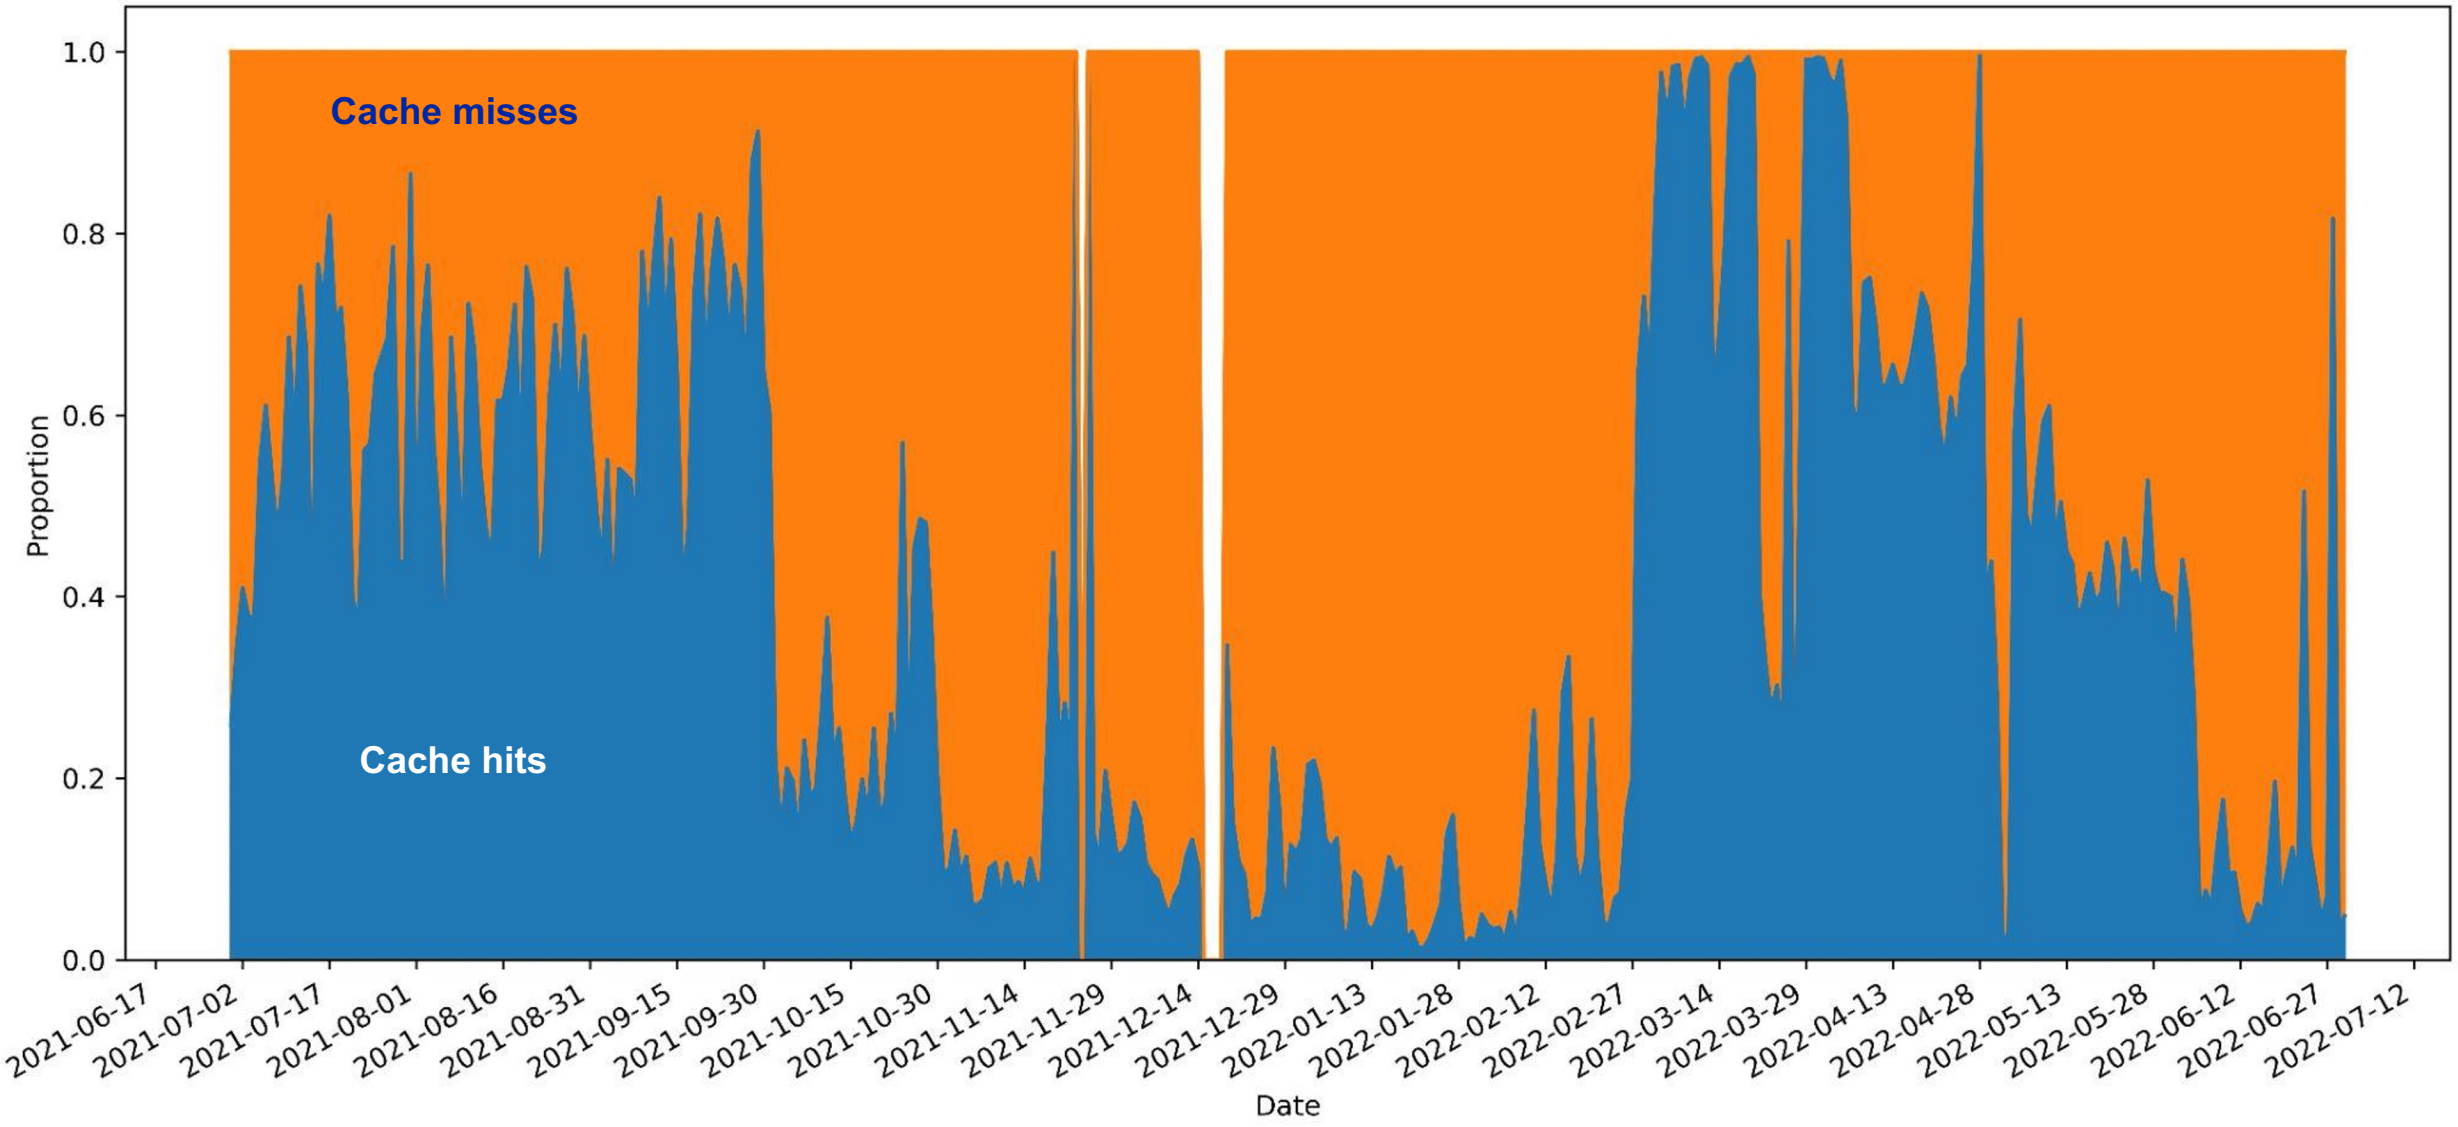

Weekly proportion of cache hits volume

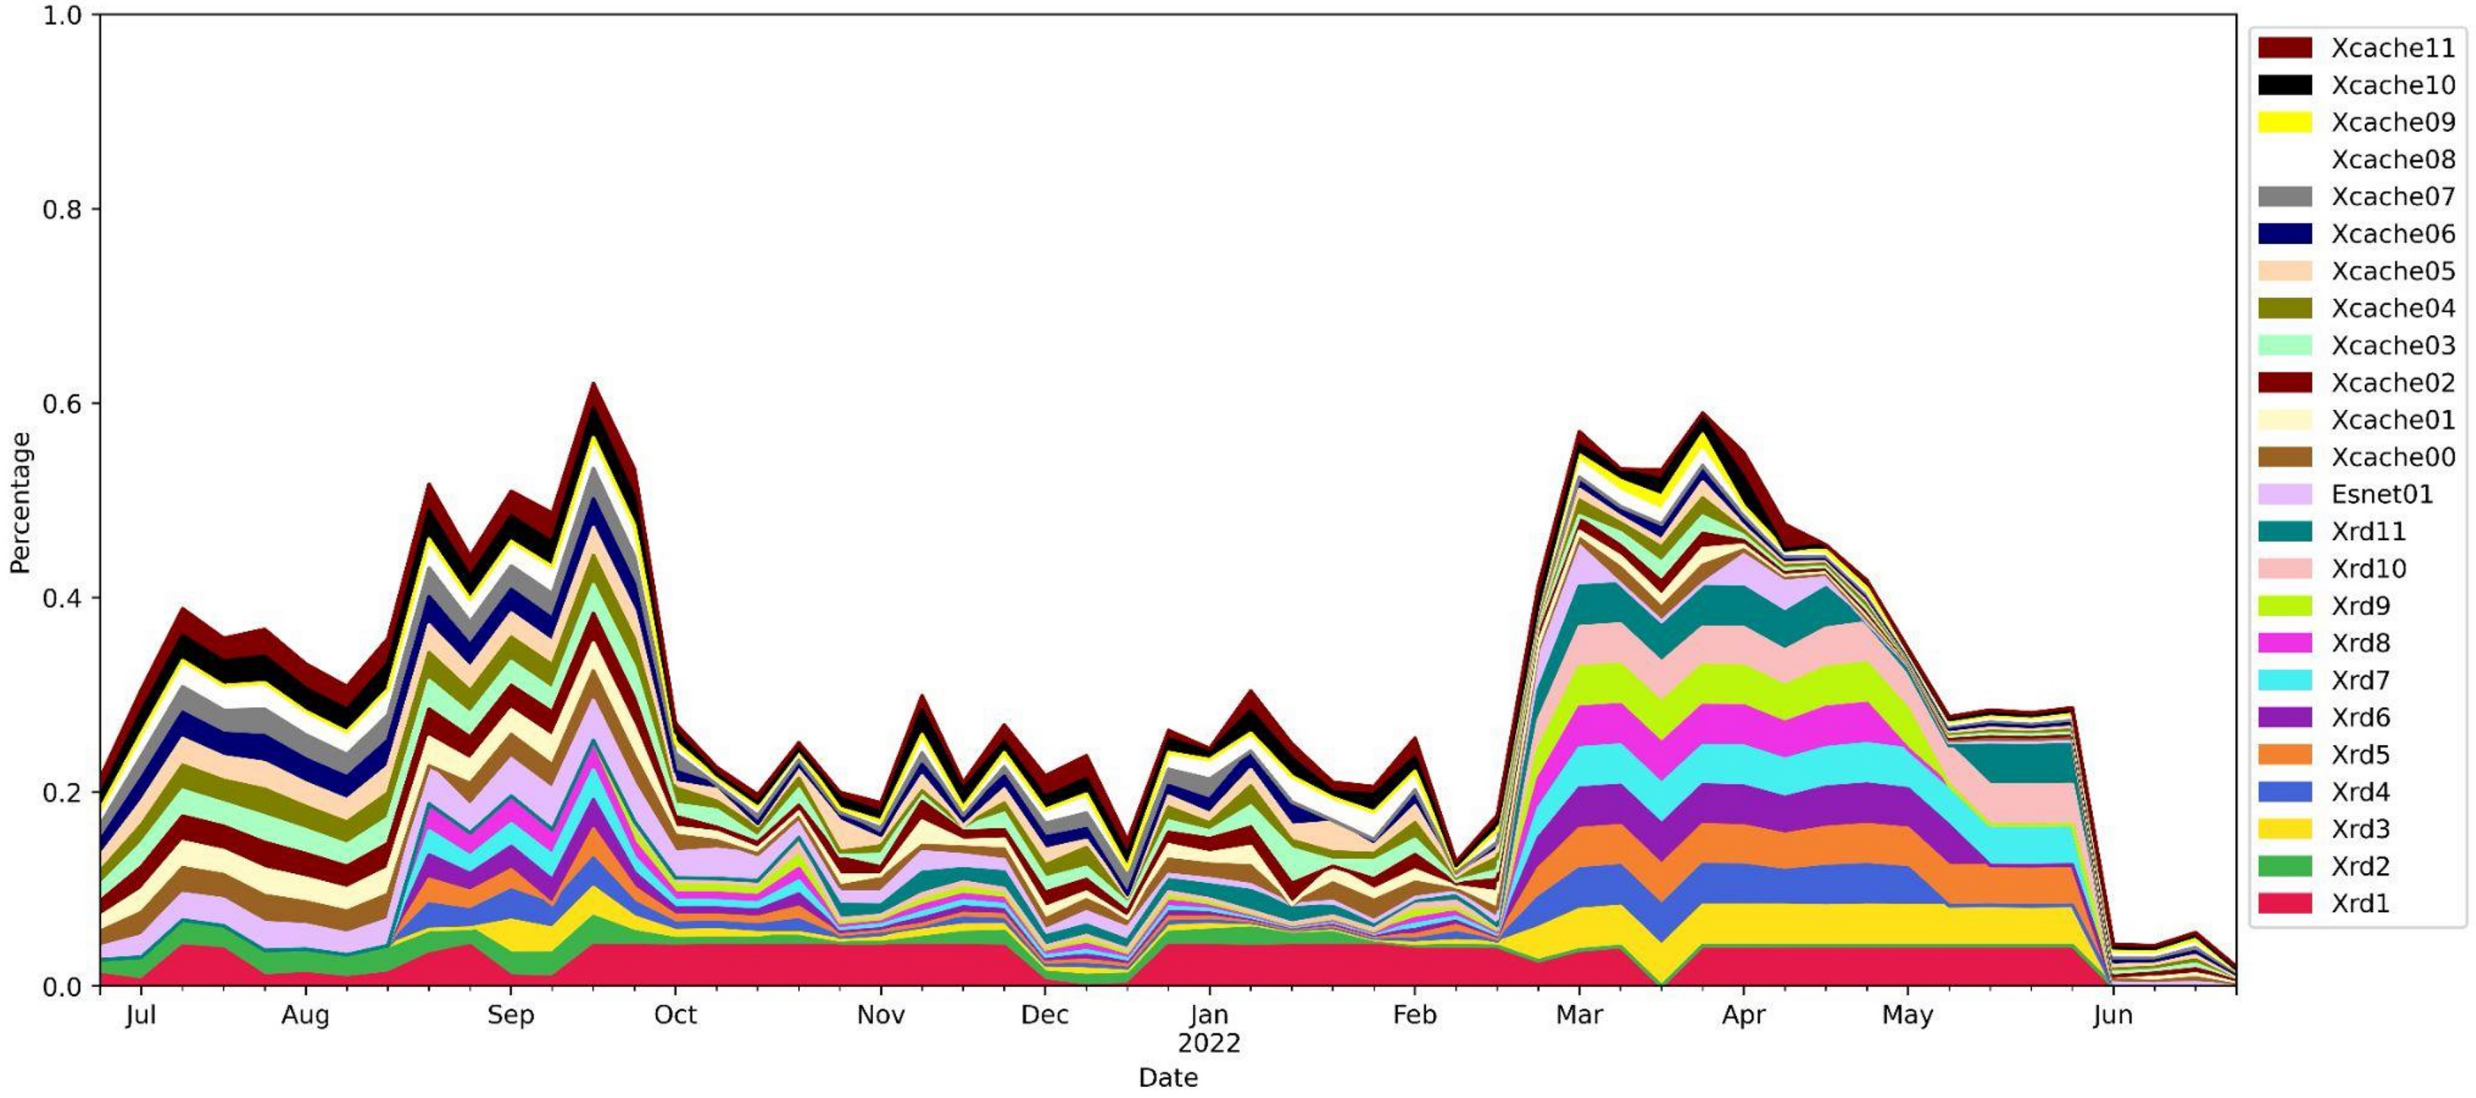

Daily network traffic demand reduction rate (traffic volume)

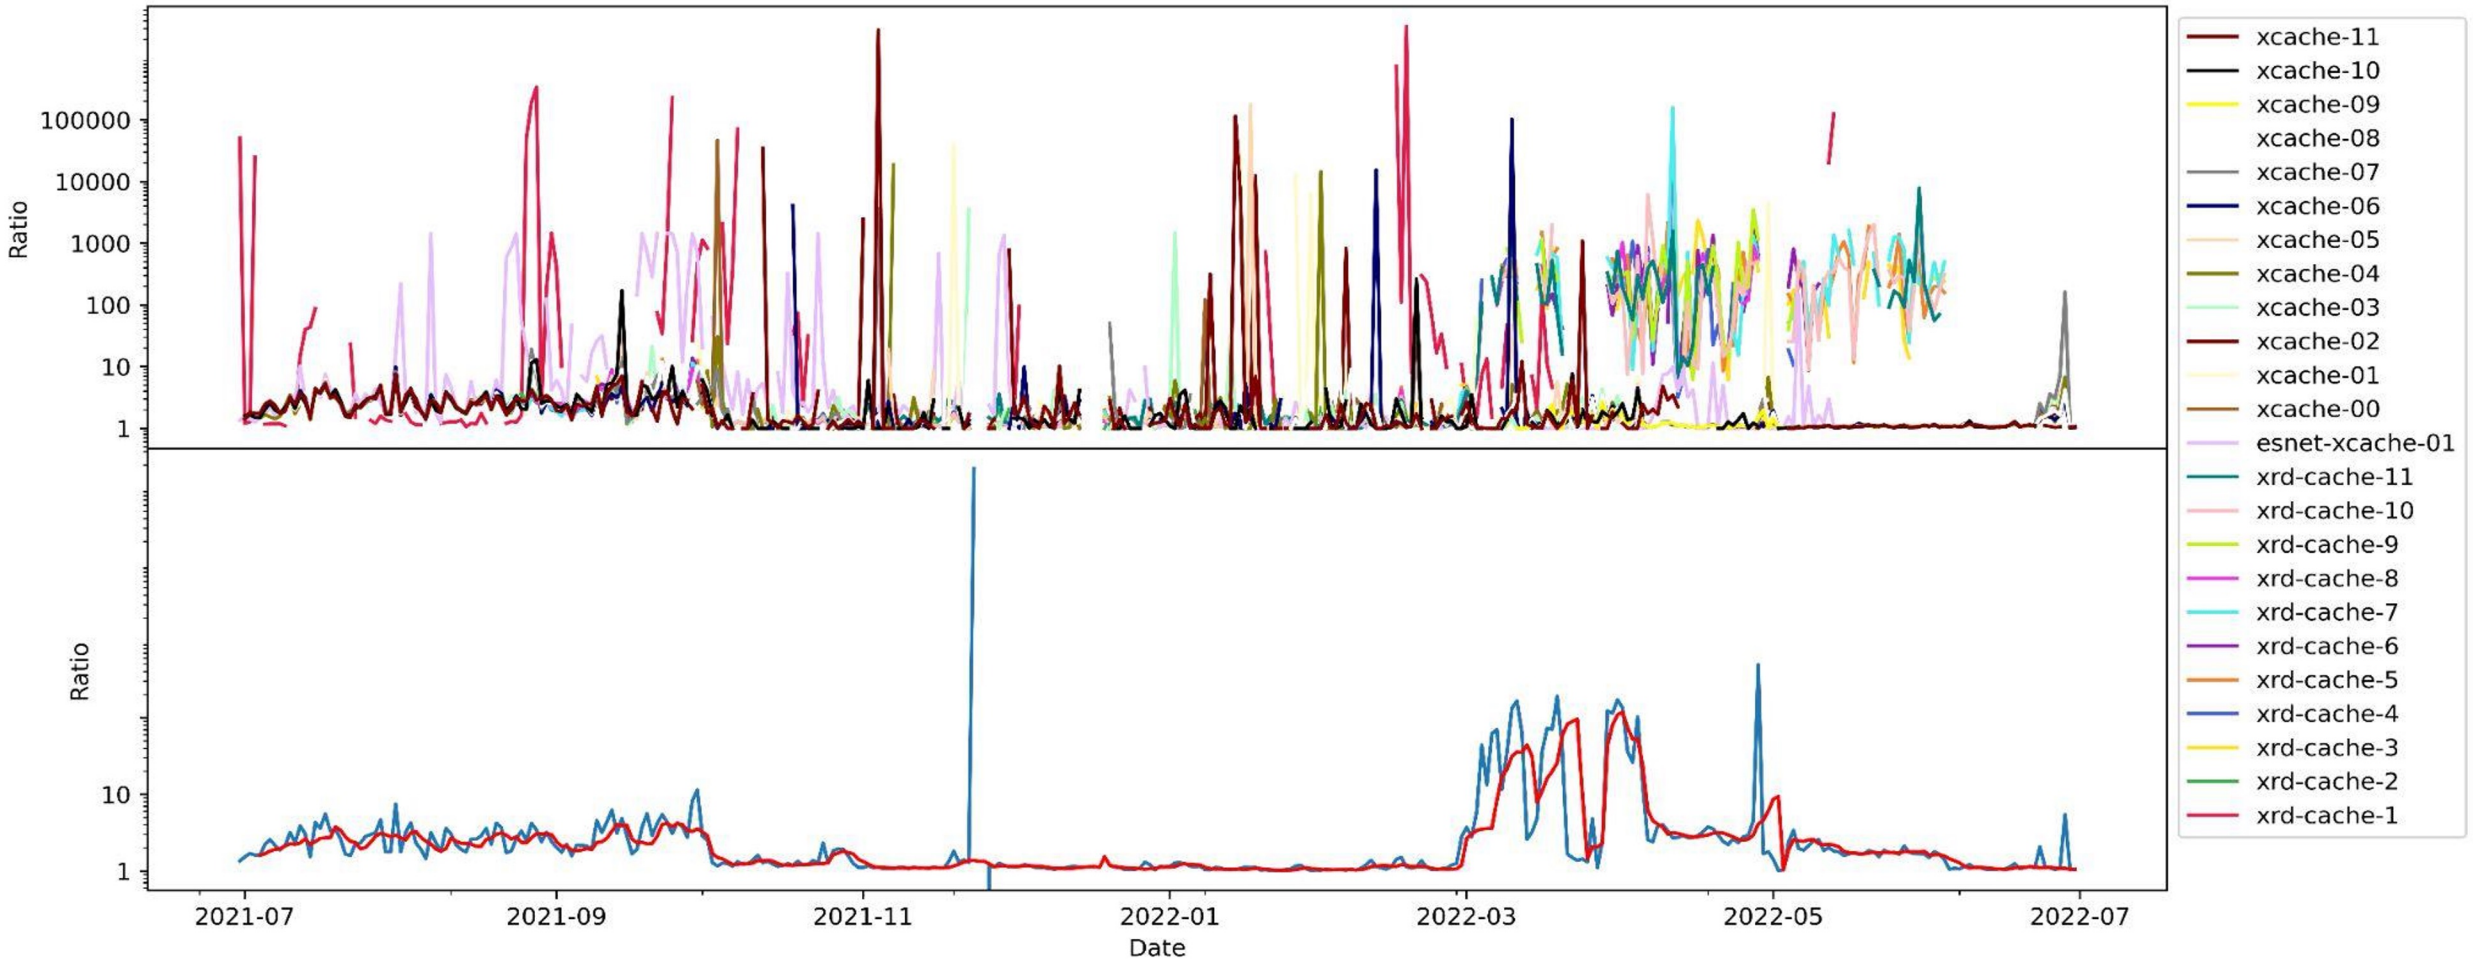

network traffic demand reduction rate = (total cache hit volume + total cache miss volume) / (total cache miss volume)

Avg traffic volume reduction for the whole period: 1.55
Avg traffic volume reduction from July 2021 - Sept 2021: 2.35

Cache utilization

- "Trivial File Catalogue" (TFC) handles user requests
  - "Local redirector" knows all caches
  - If cache hit, local redirector routes to the node
  - if cache miss, an XRootD client to fetch the file from the national XRootD data federation
- Handling cache miss
  - Cache miss file goes to nodes with empty space first
    - Nodes with empty spaces receive most of data accesses
      - cache miss goes to empty space
      - cache hits mostly on newly transferred data, which are previous cache misses
    - When new nodes added
      - Caltech nodes (Xrd 3 - 8, 11) added around Aug 26th, 2021
      - new Caltech nodes (Xrd 9 - 10) added around Sept 30th, 2021
  - Delete old data if space are full
    - After the system running for sufficient time without adding new nodes
    - All cache nodes divide the data access more evenly

Daily cache hit volume

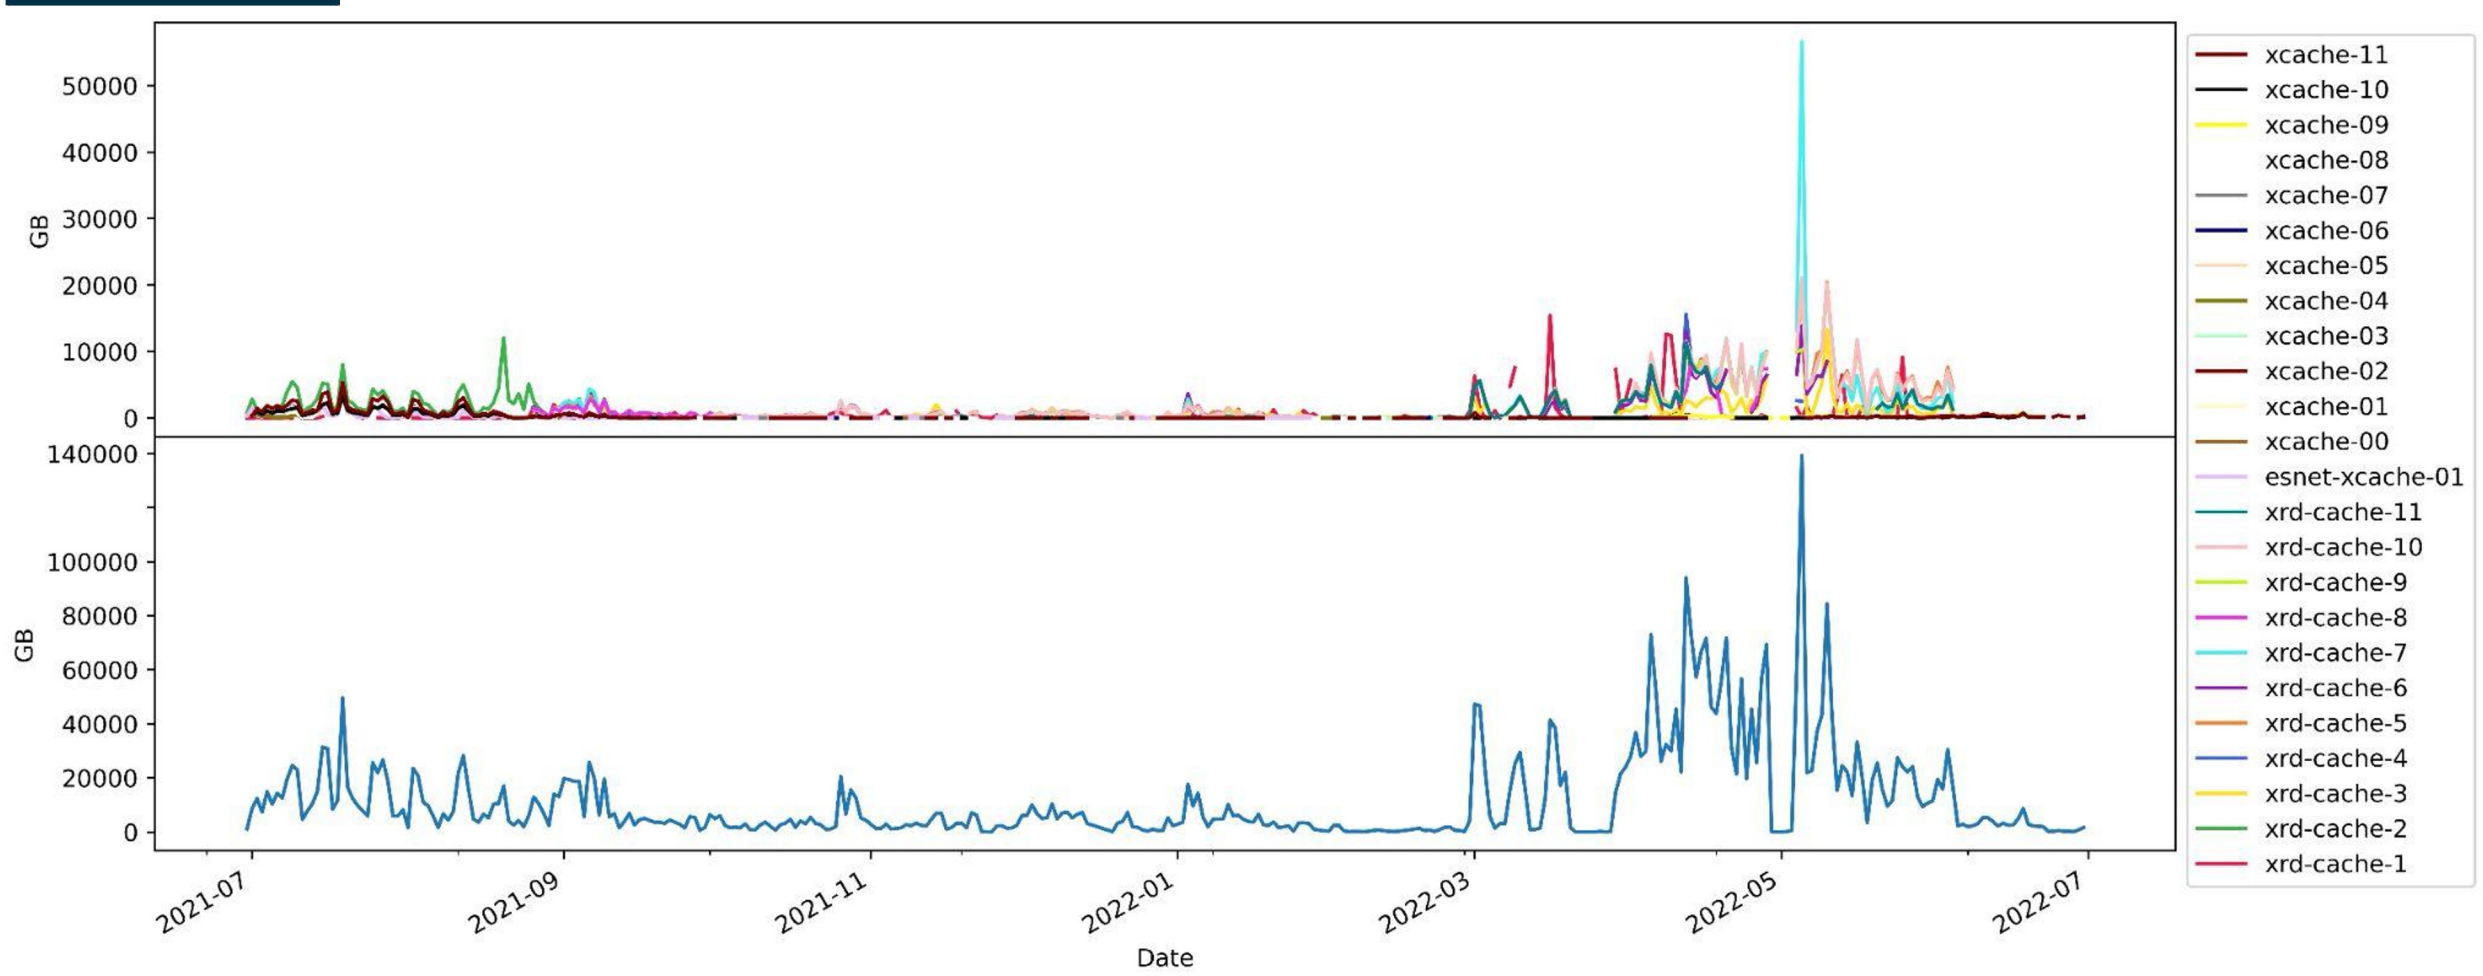

Daily cache hit volume with 2-day moving average

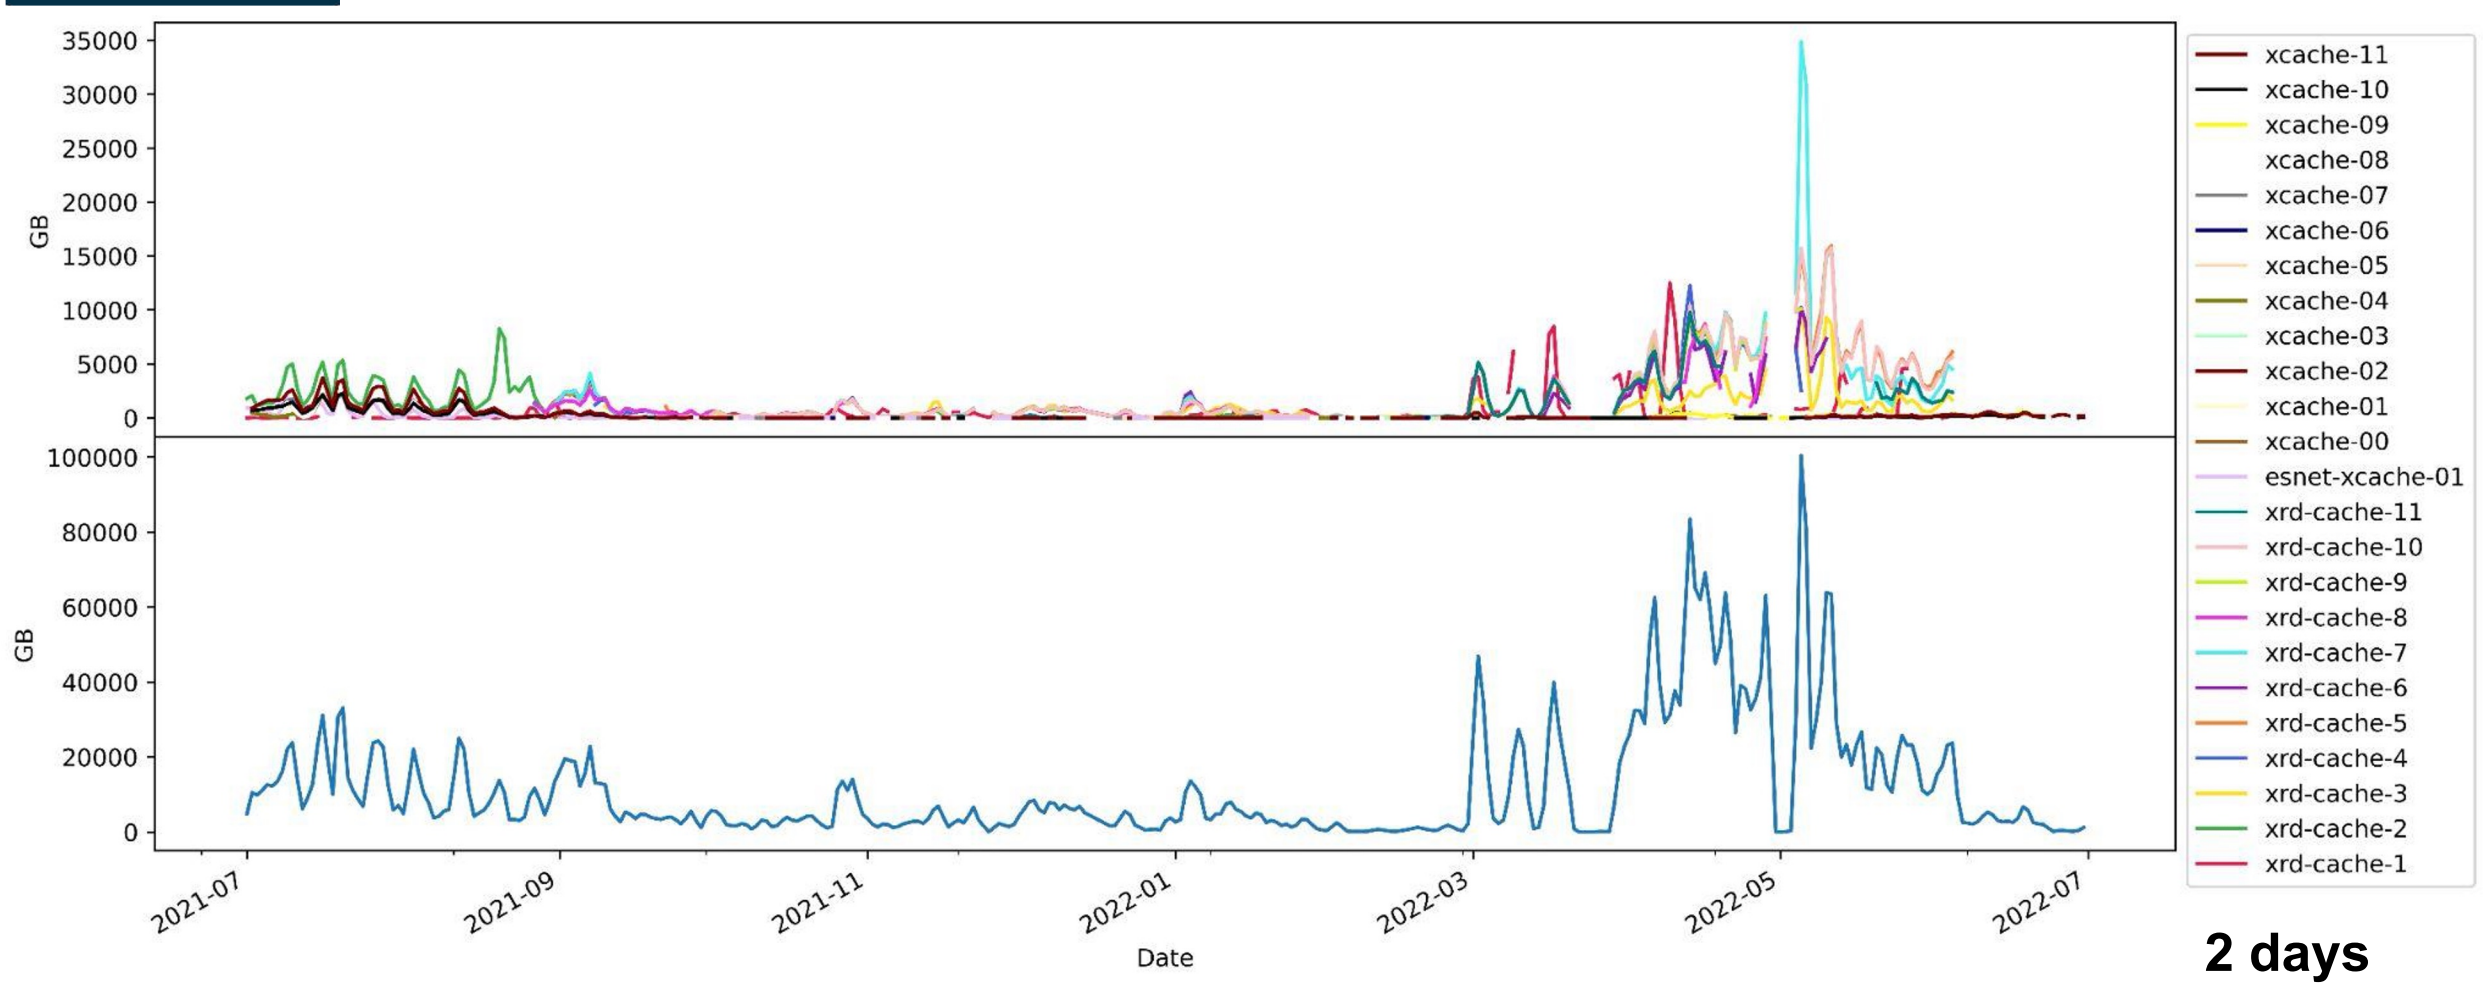

2 days

Daily cache hit volume with 3-day moving average

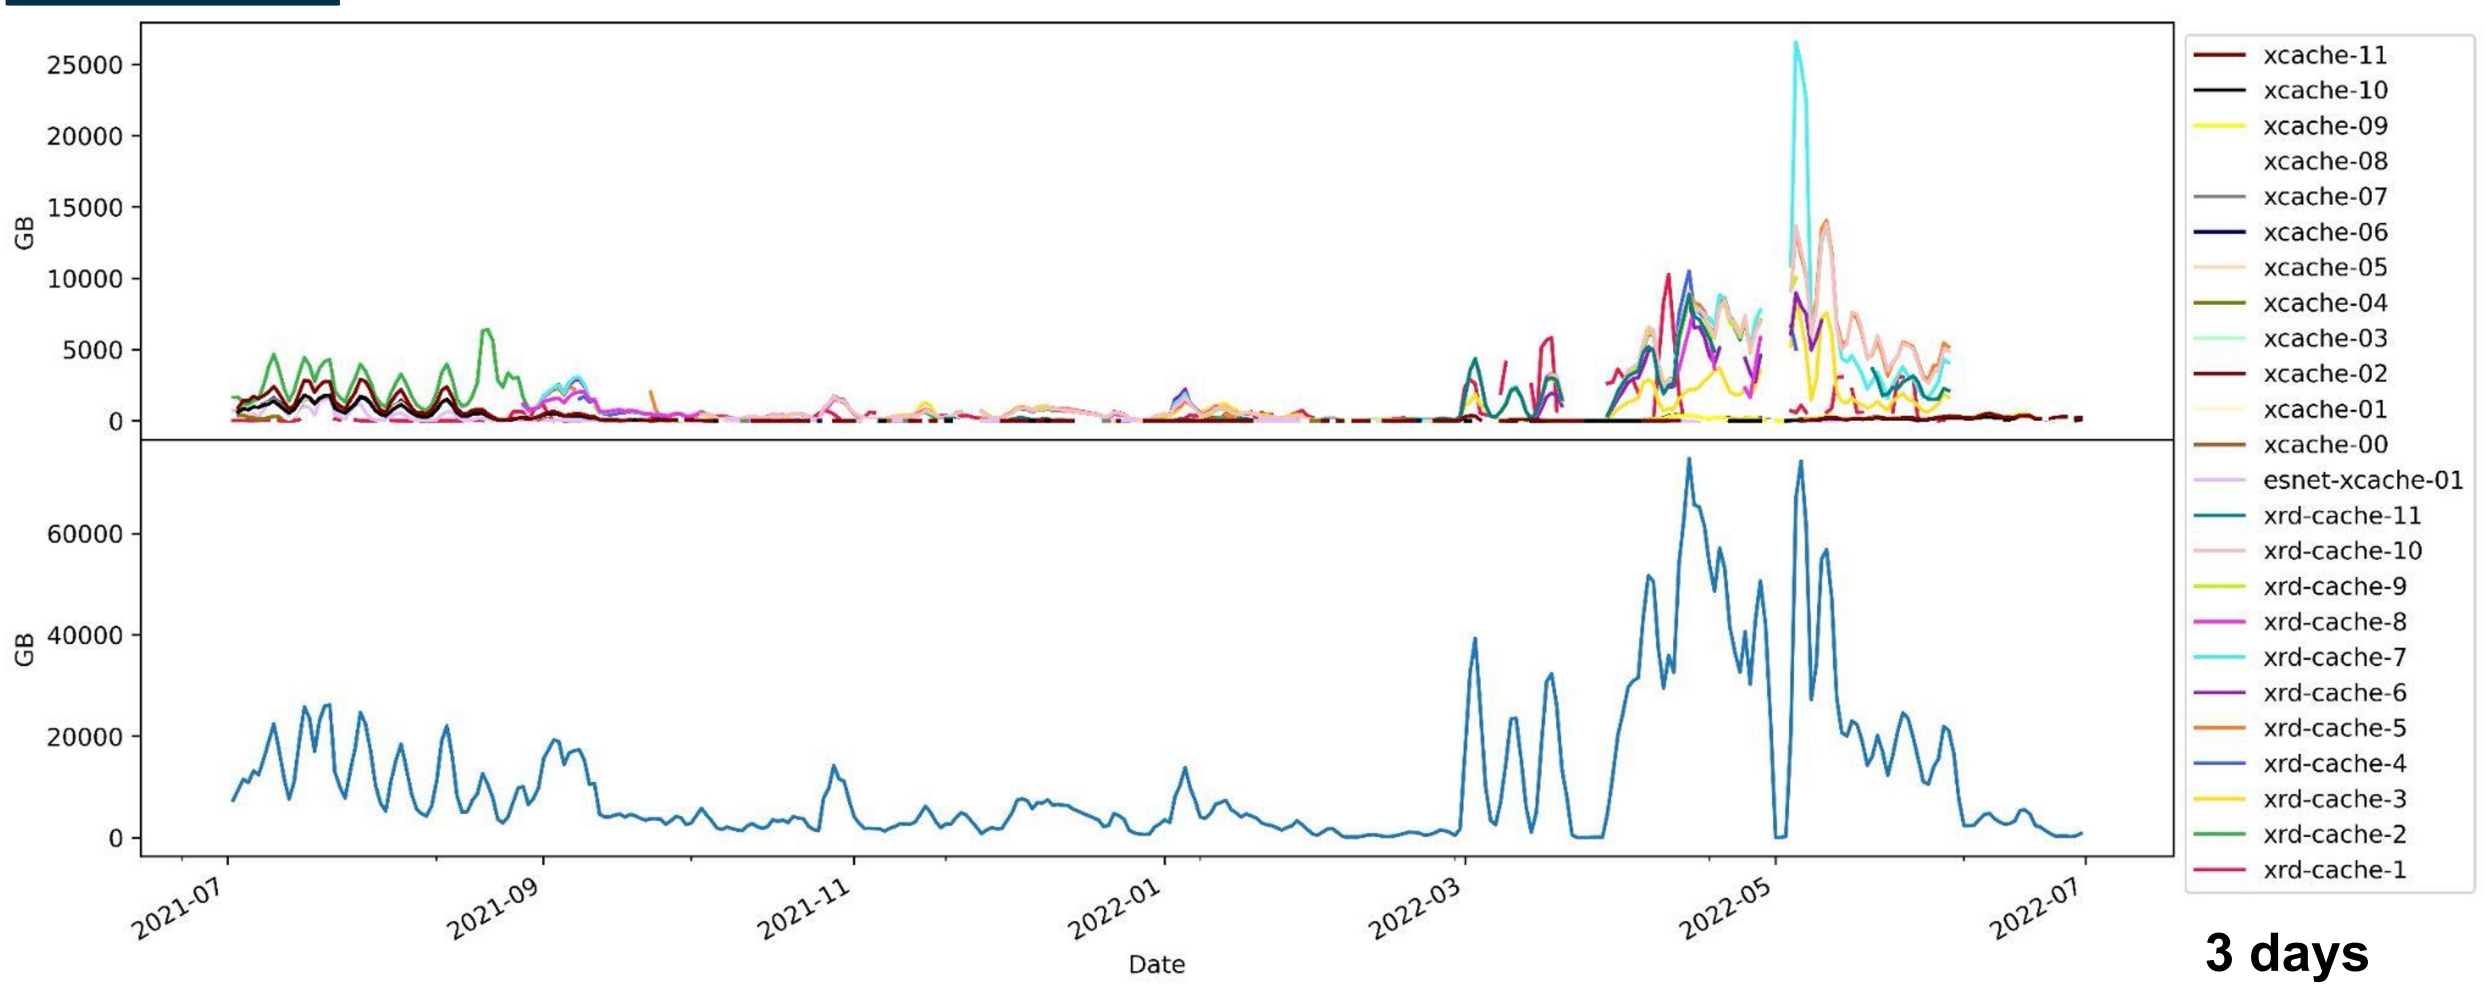

3 days

Daily cache hit volume with 5-day moving average

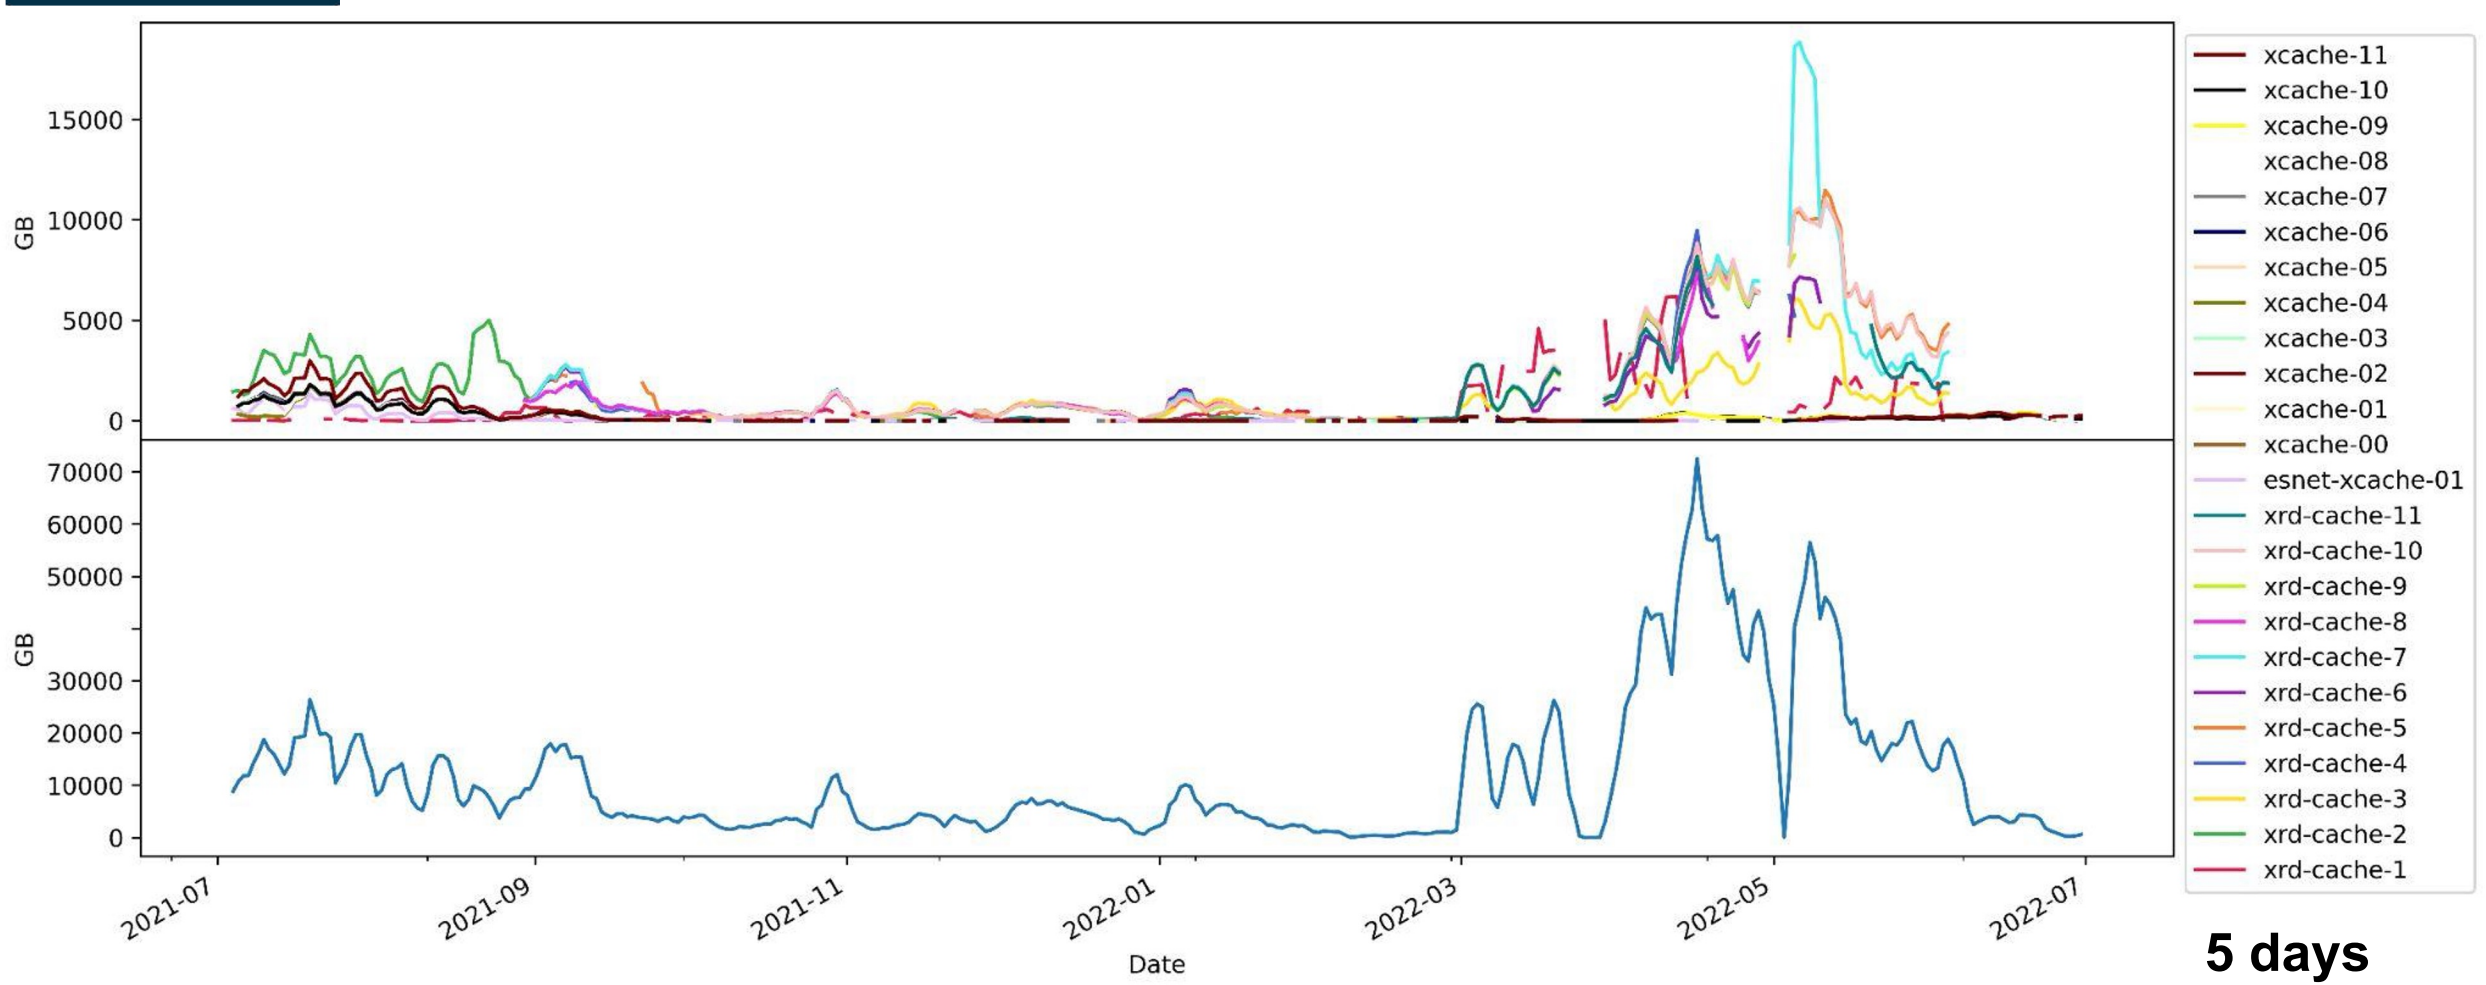

5 days

Daily cache hit volume with 7-day moving average

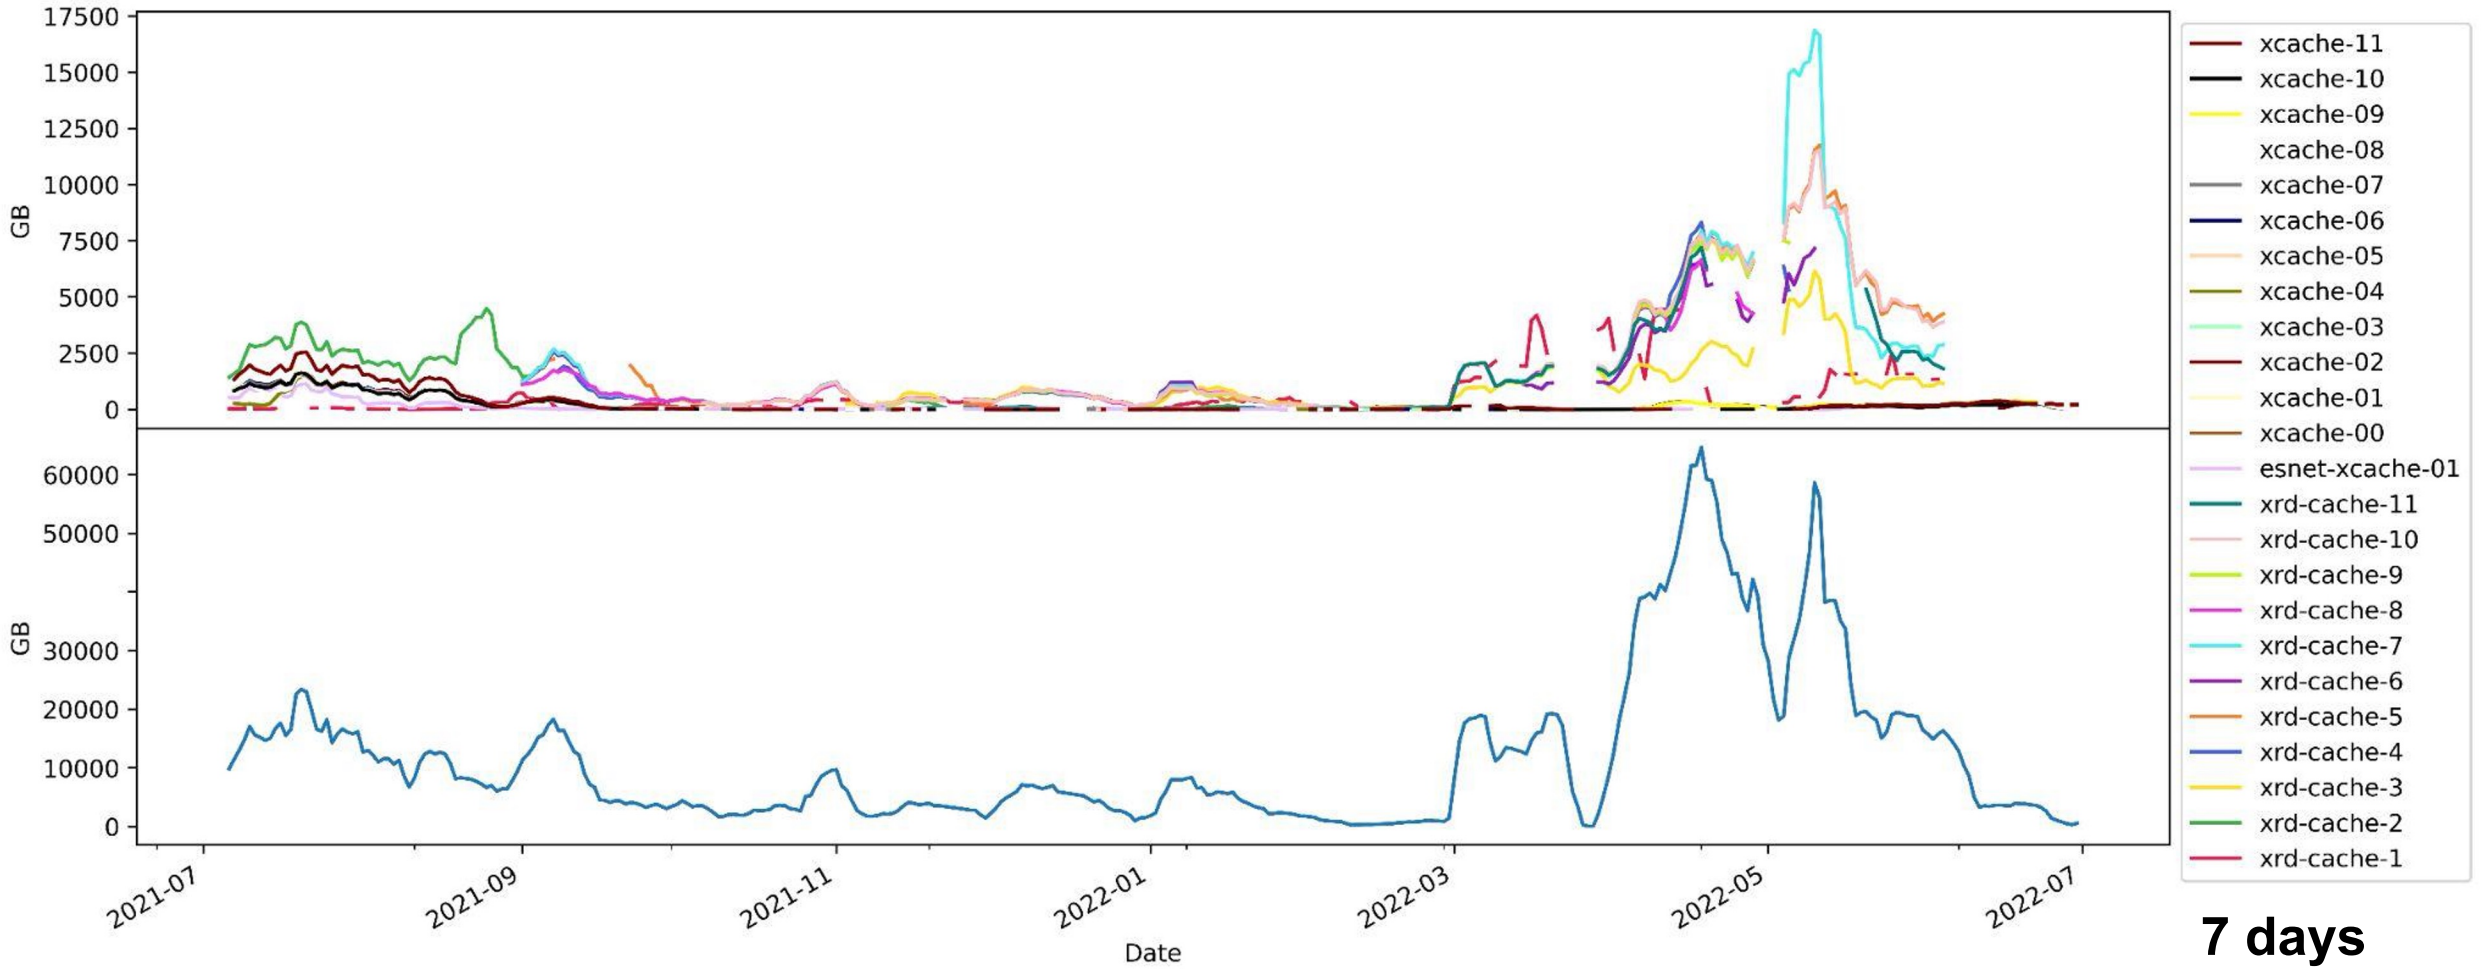

7 days

Daily cache hit volume with 10-day moving average

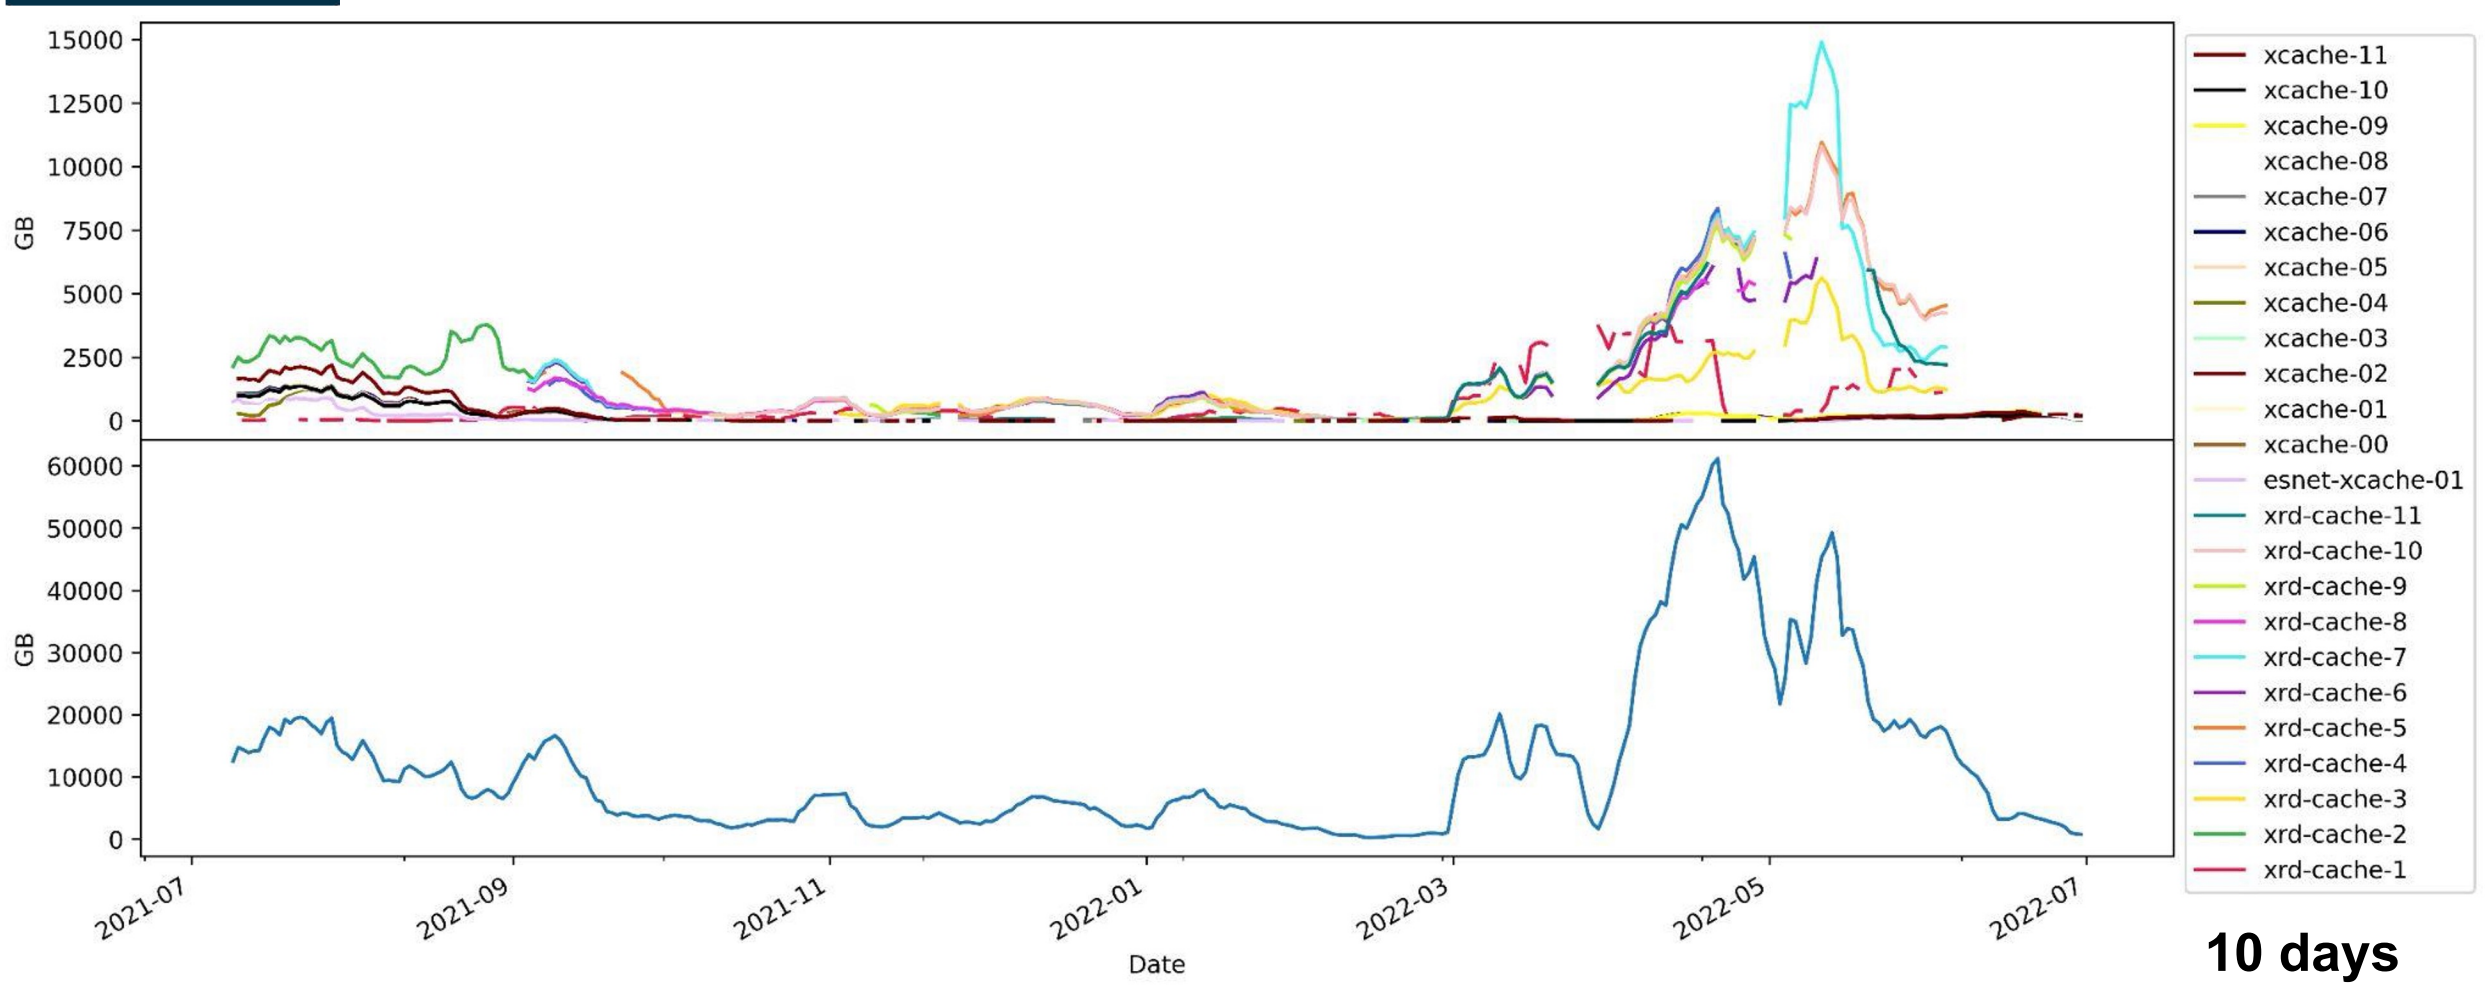

10 days

Daily cache hit volume with 14-day moving average

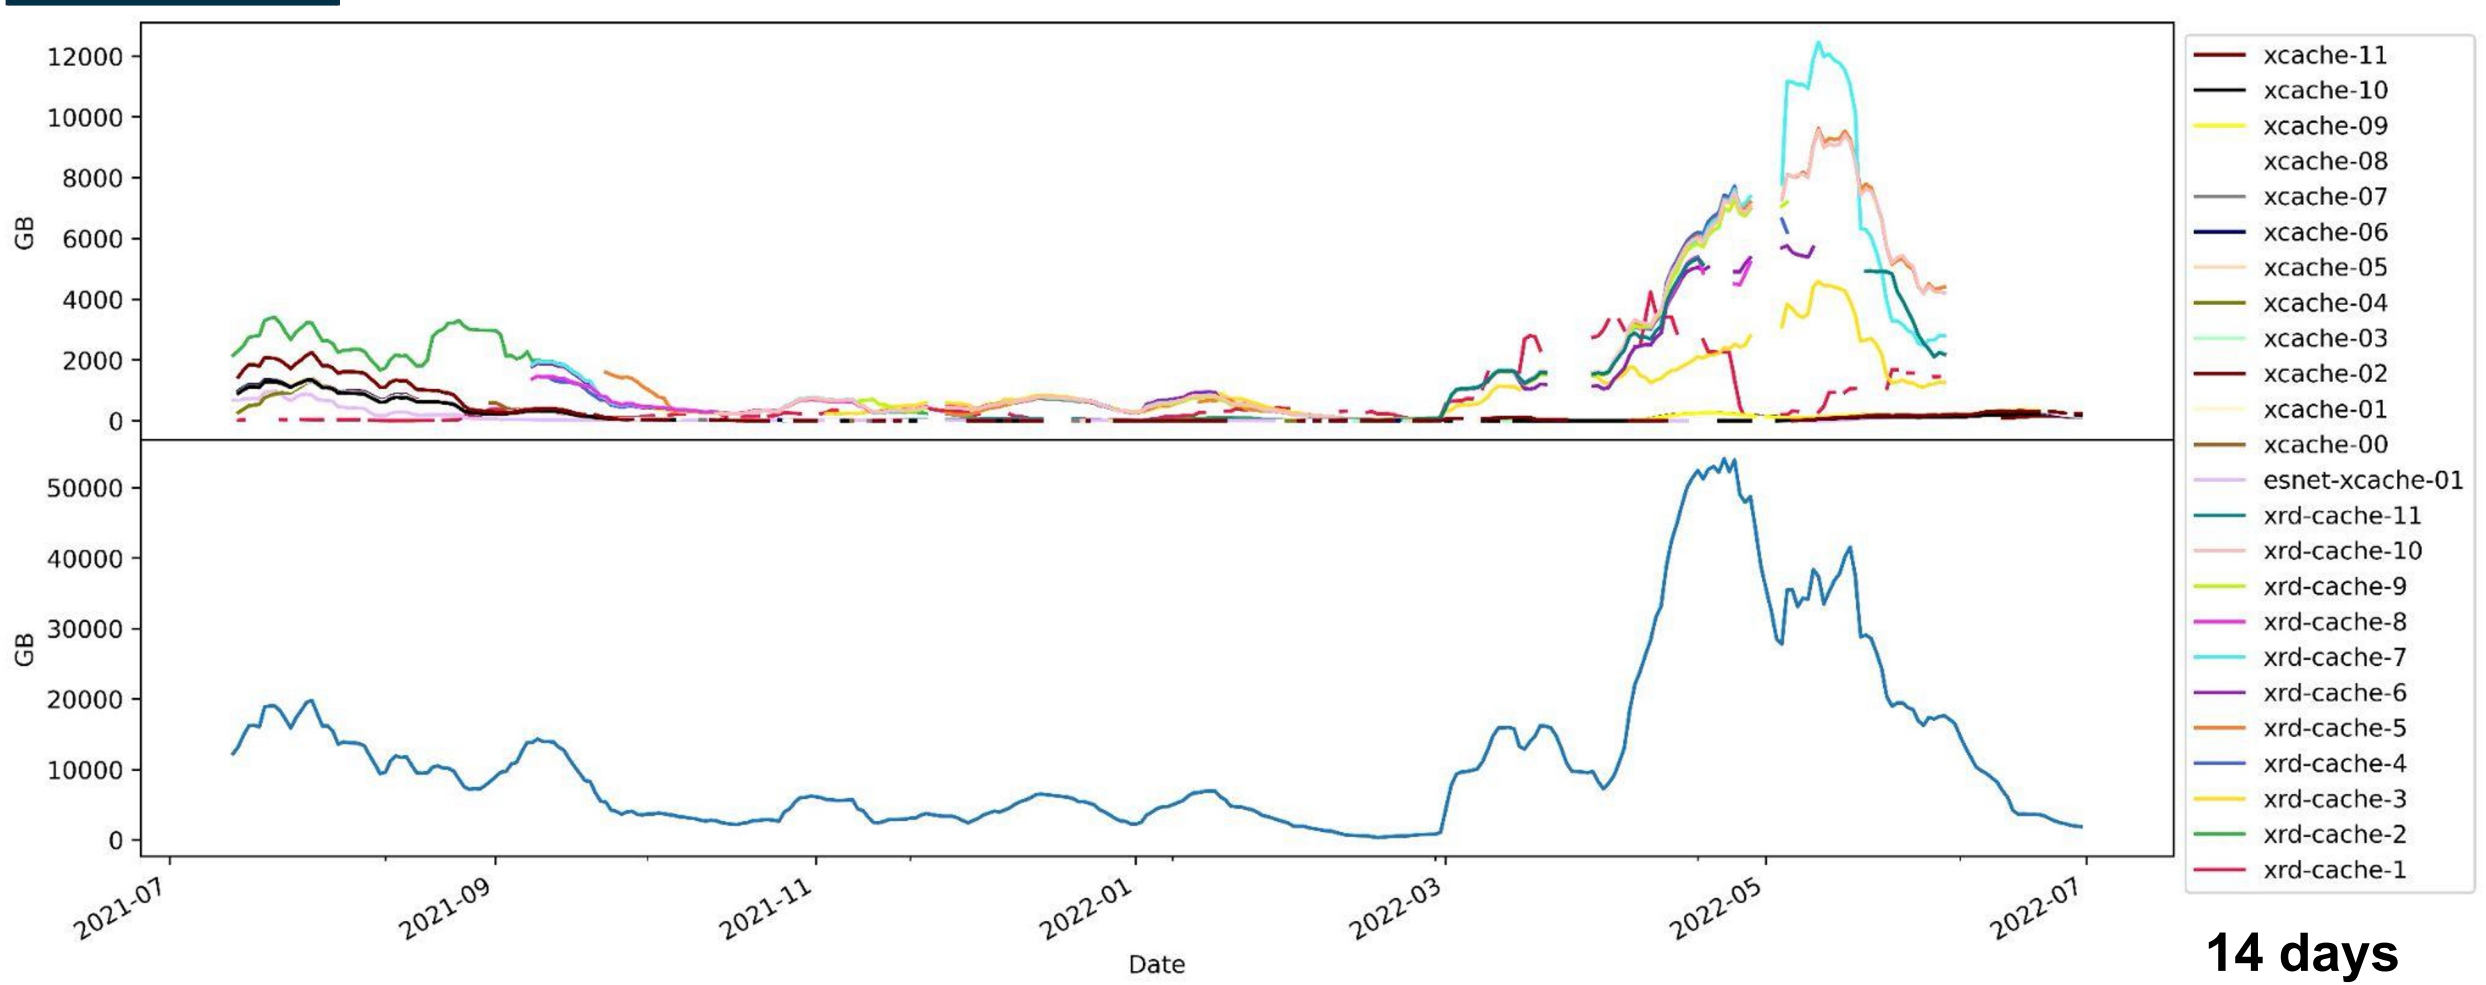

14 days

Daily cache hit volume with 21-day moving average

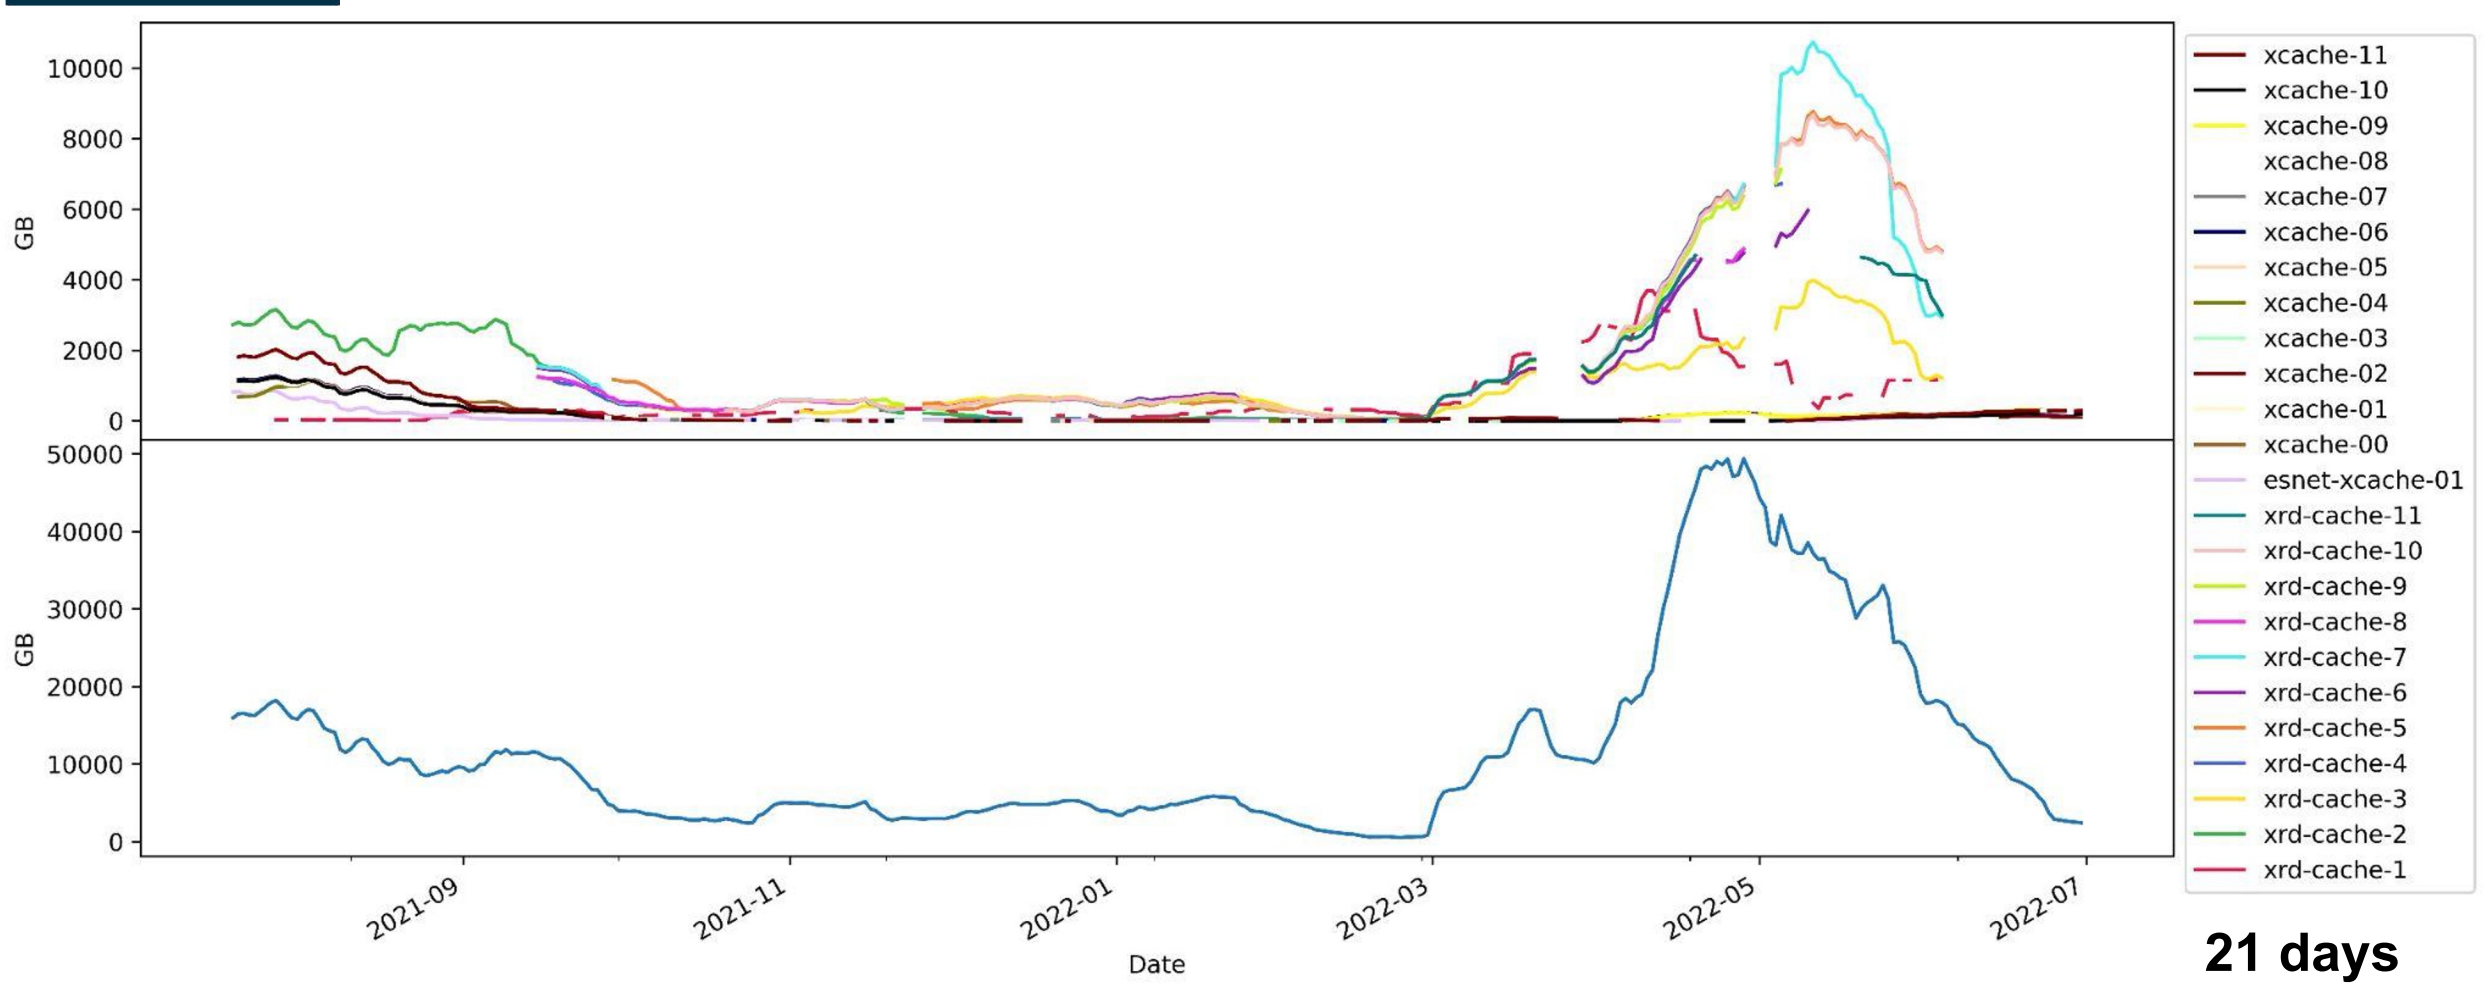

21 days

Daily cache hit volume with 28-day moving average

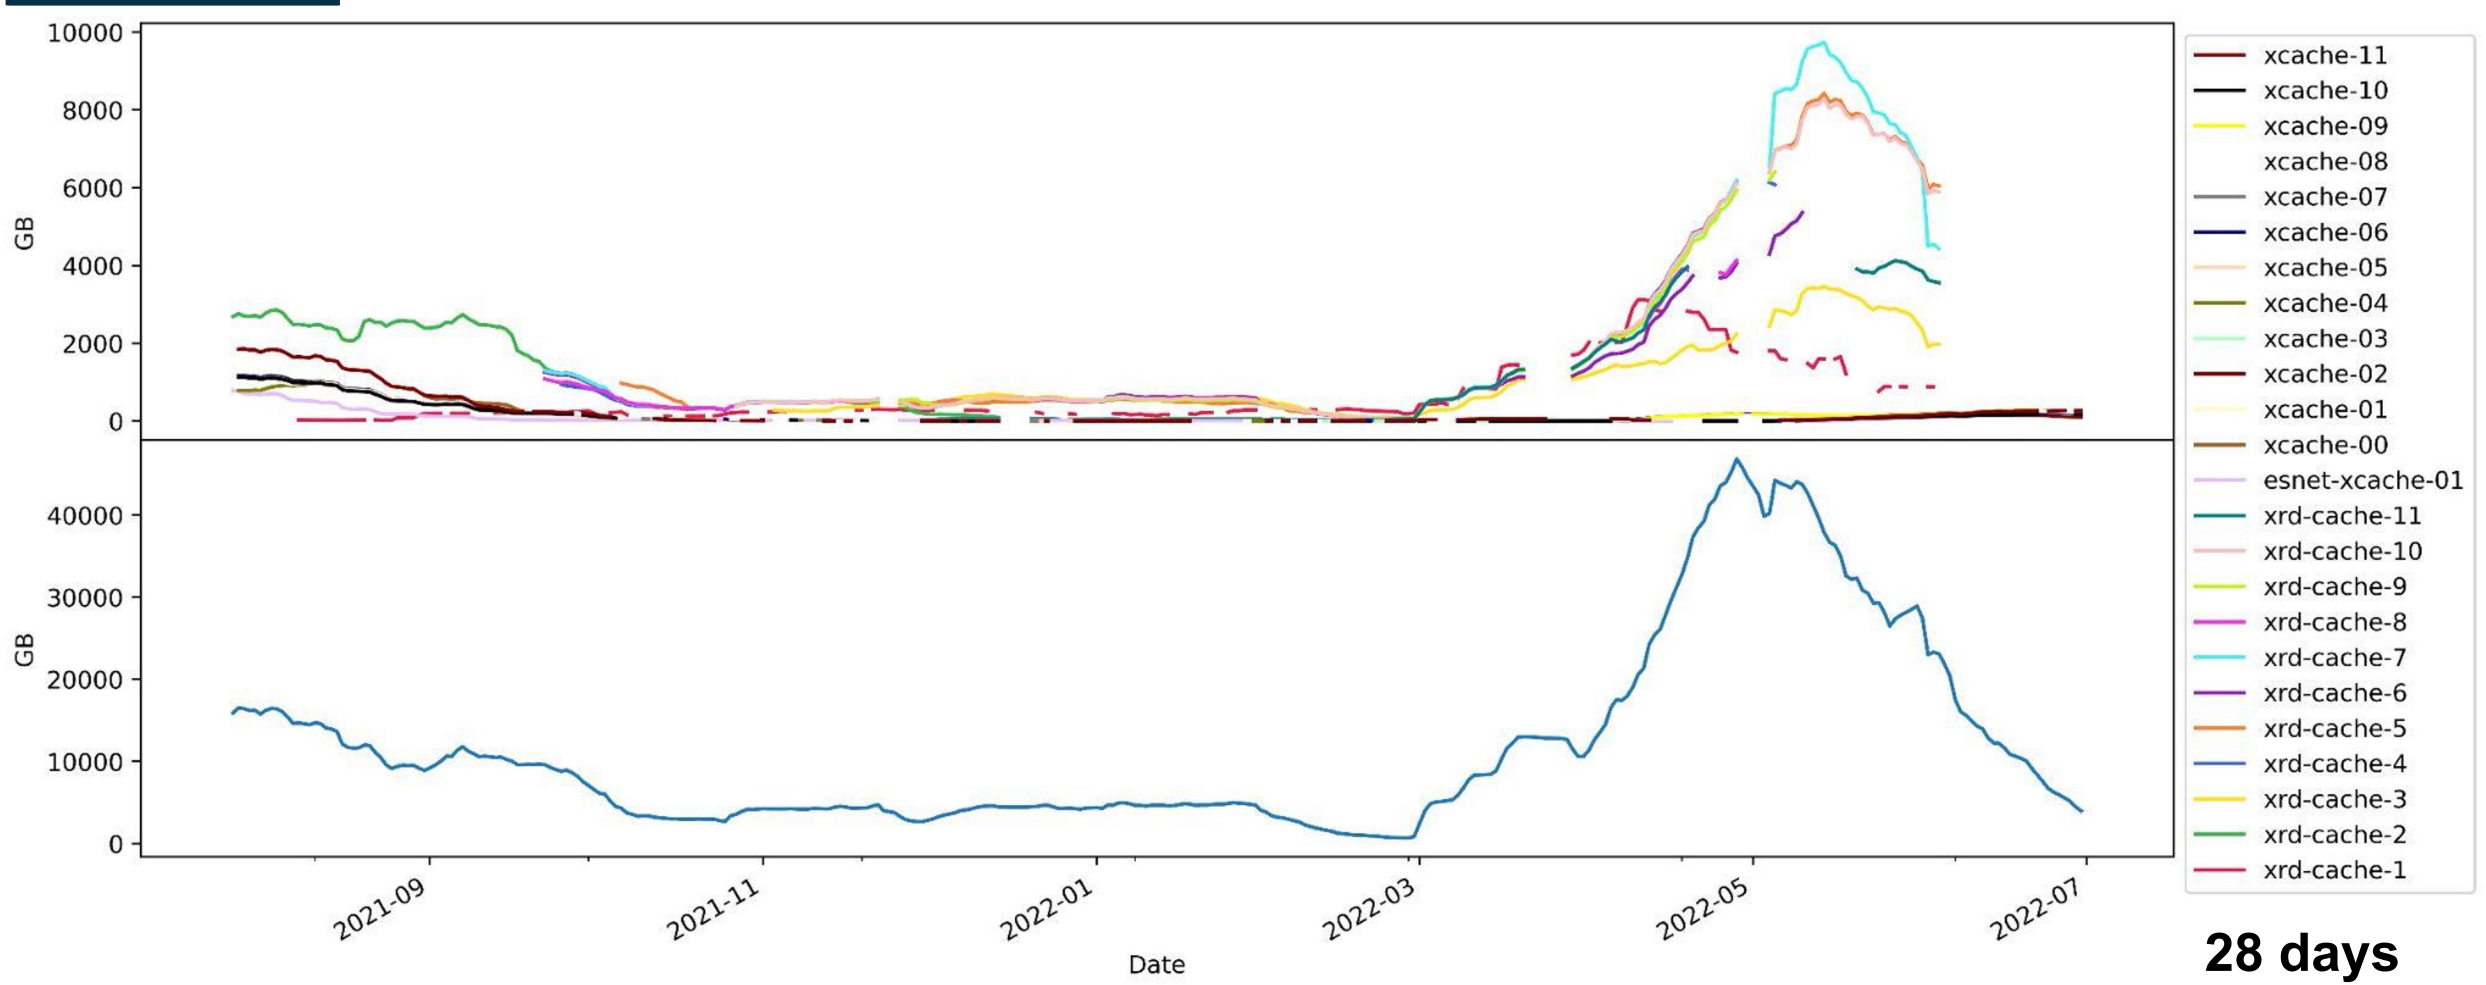

28 days

Daily cache miss volume

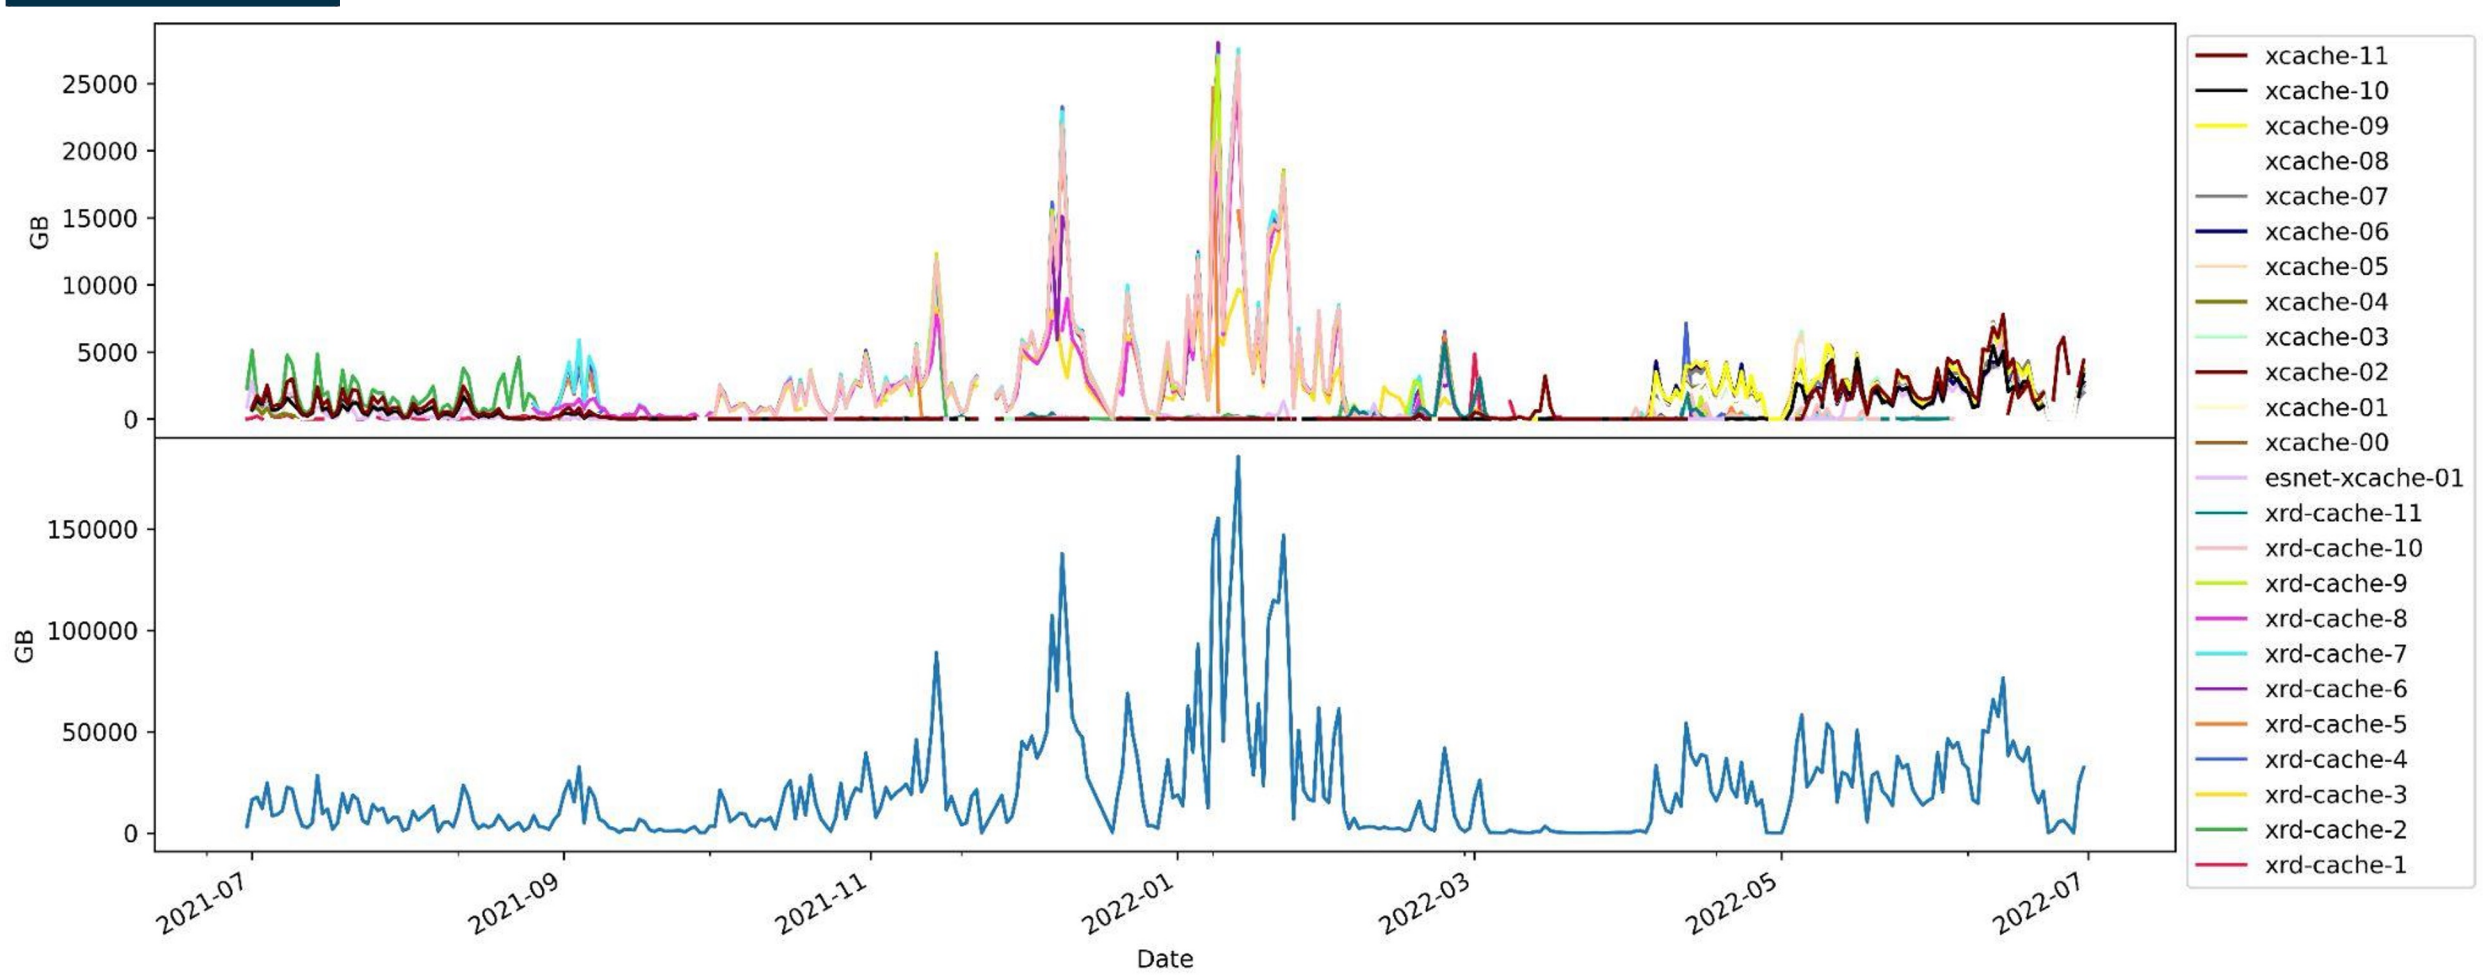

Daily cache miss volume with 2-day moving average

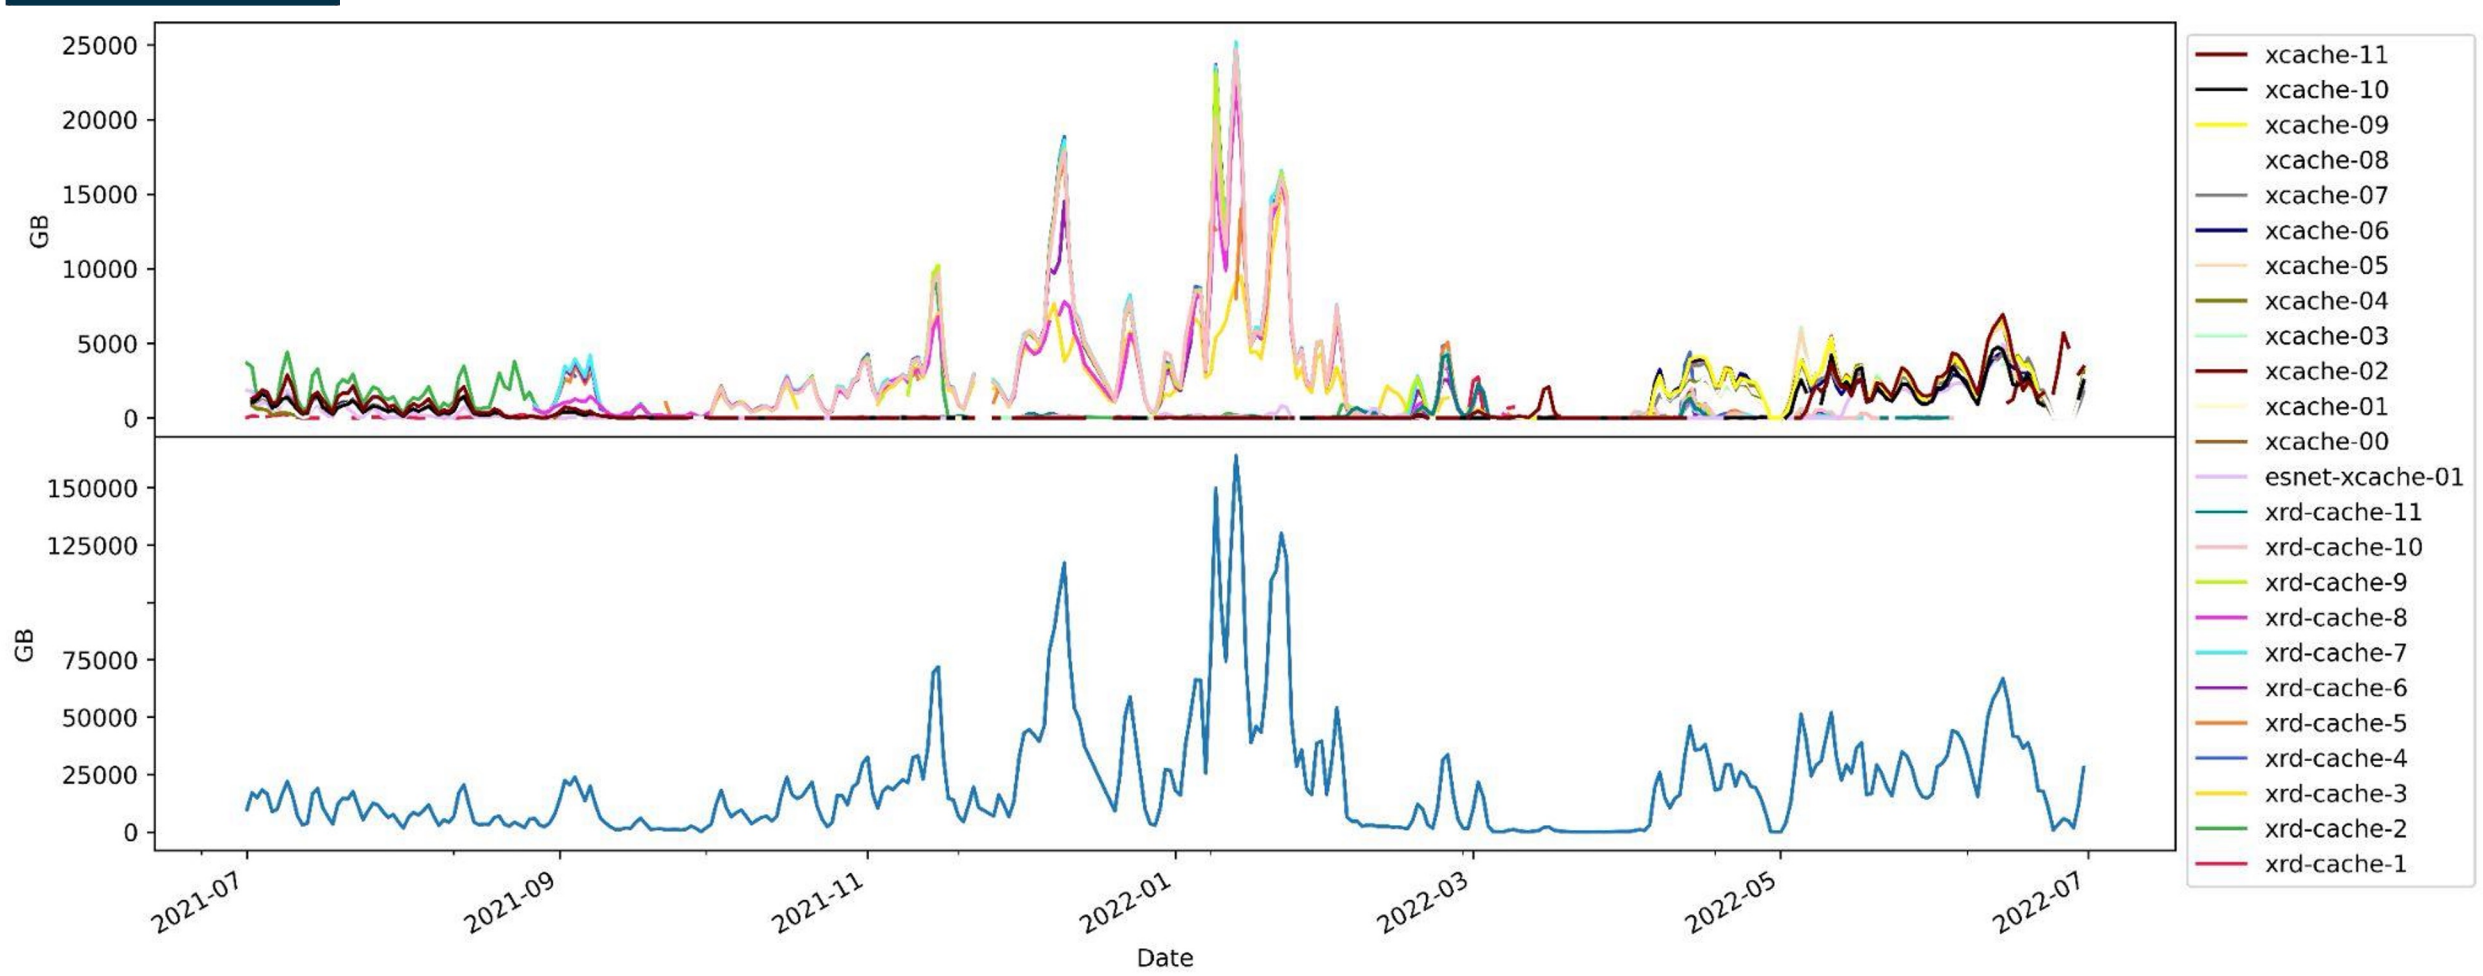

Daily cache miss volume with 3-day moving average

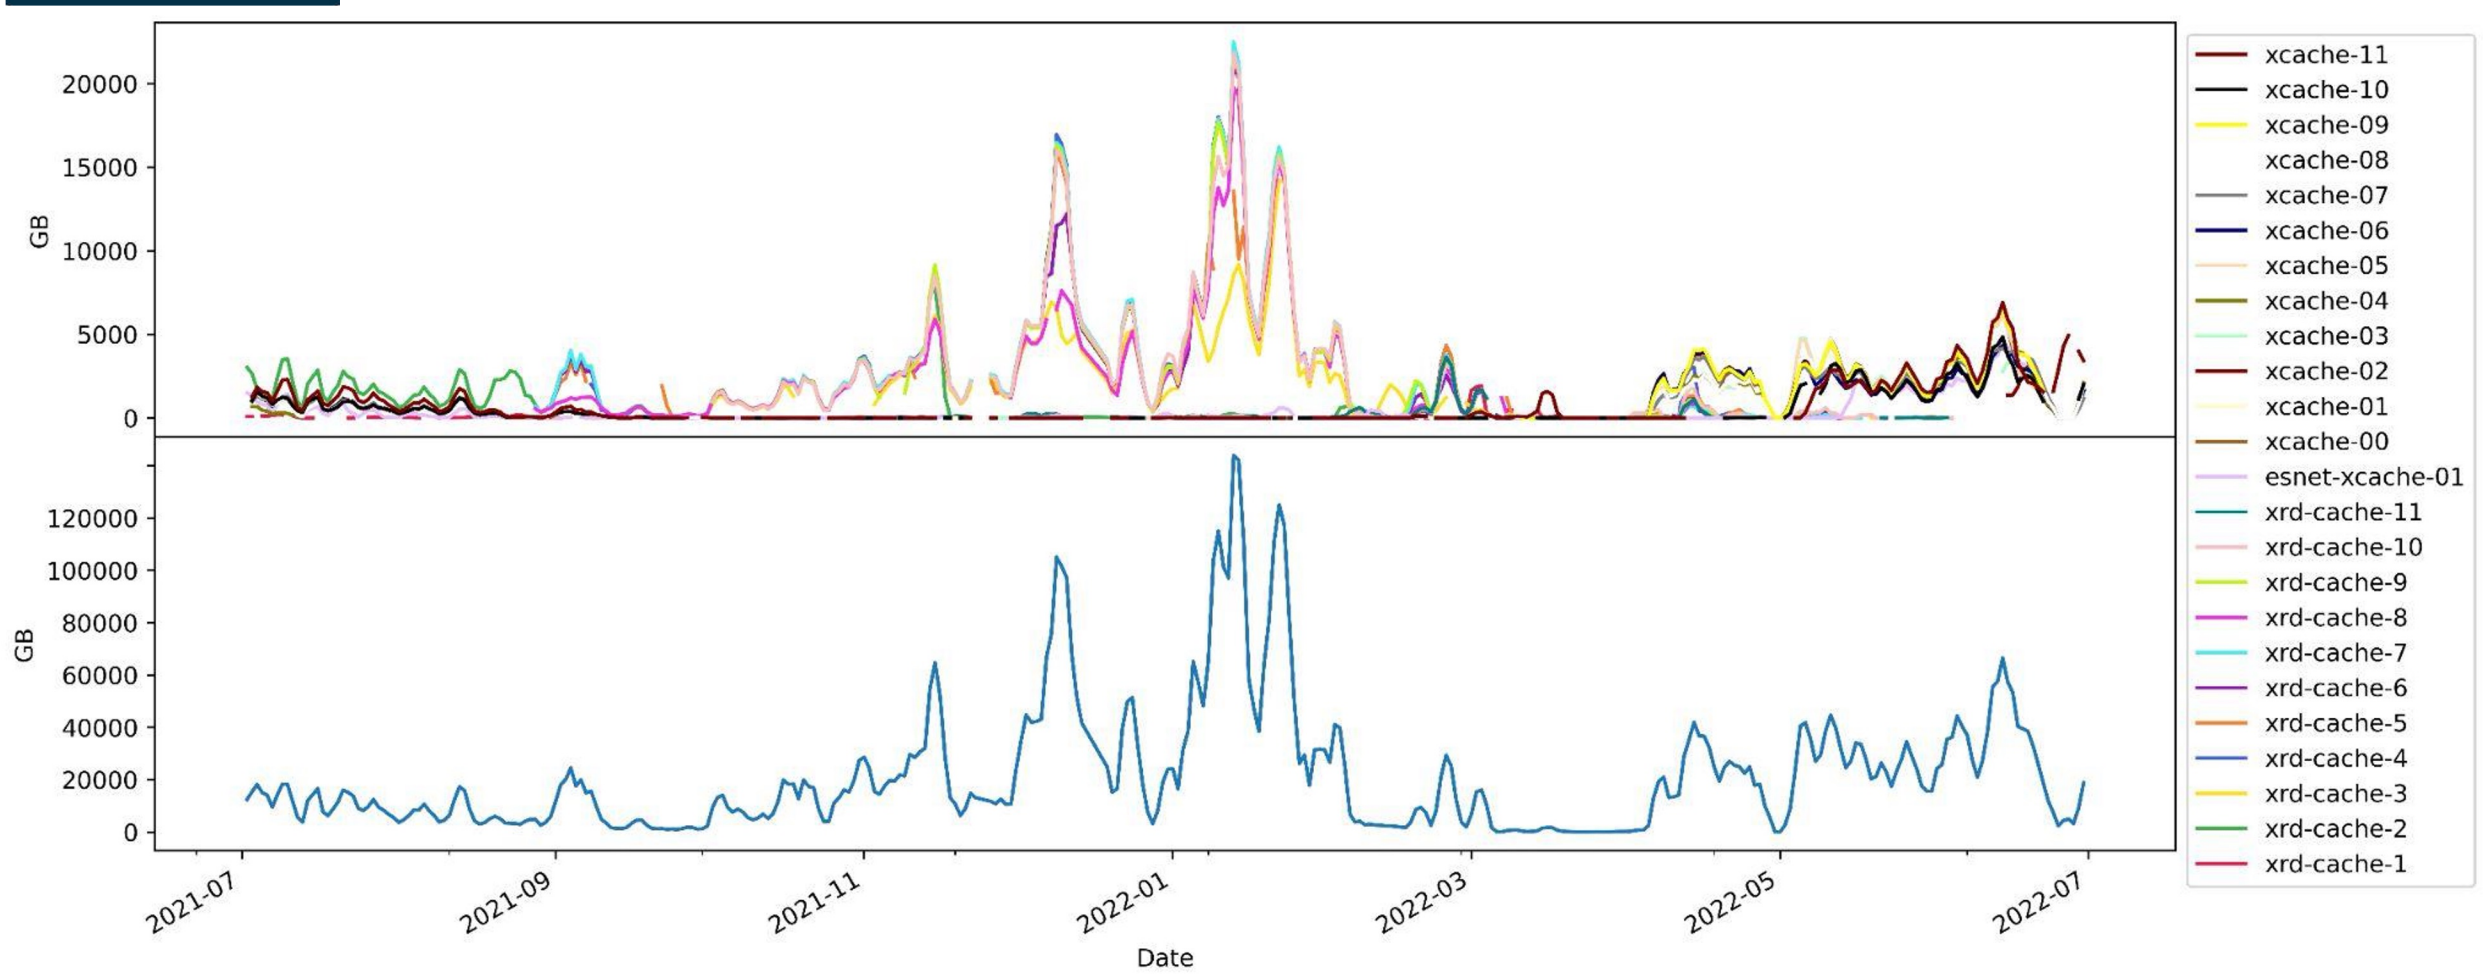

Daily cache miss volume with 5-day moving average

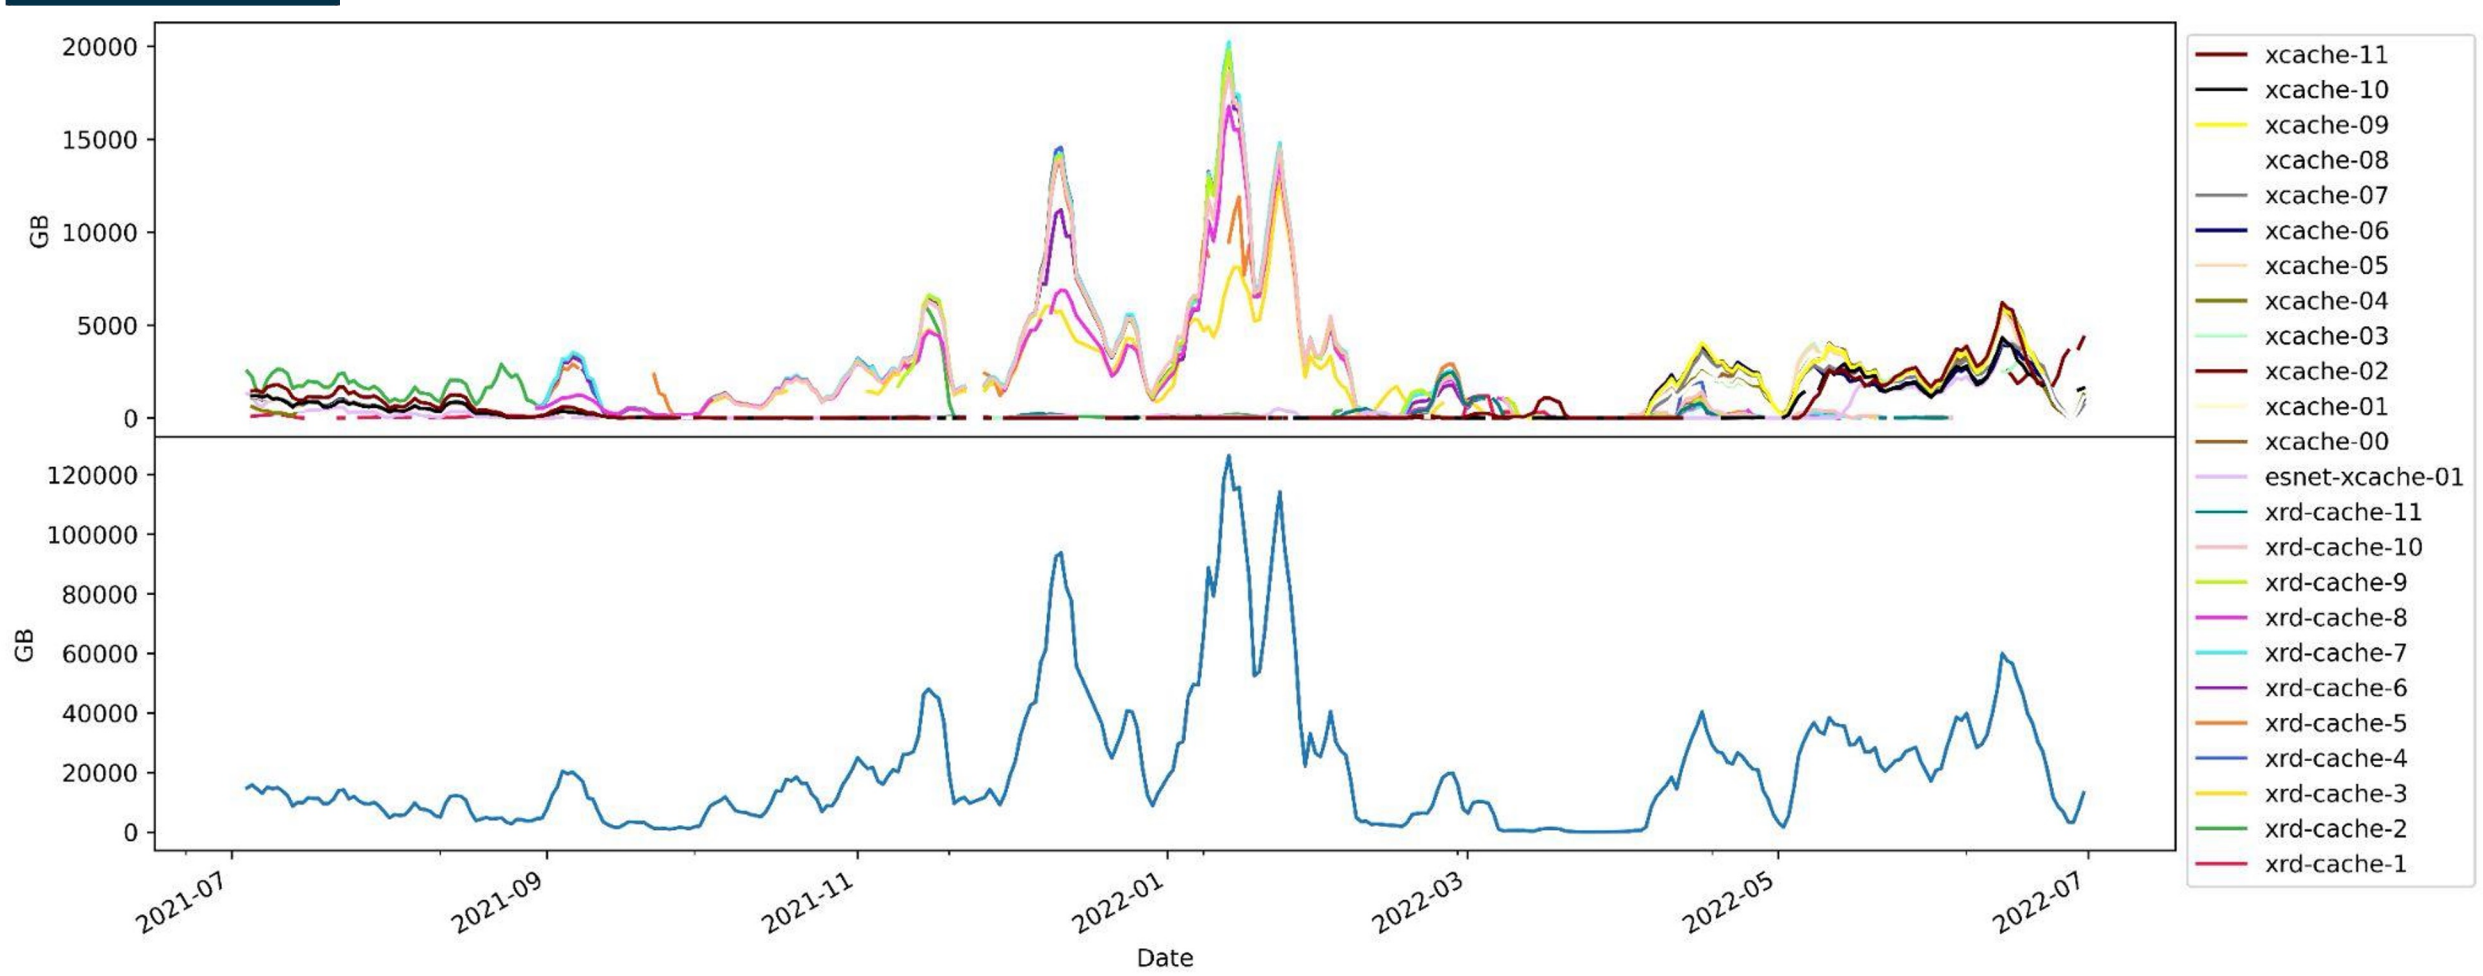

Daily cache miss volume with 7-day moving average

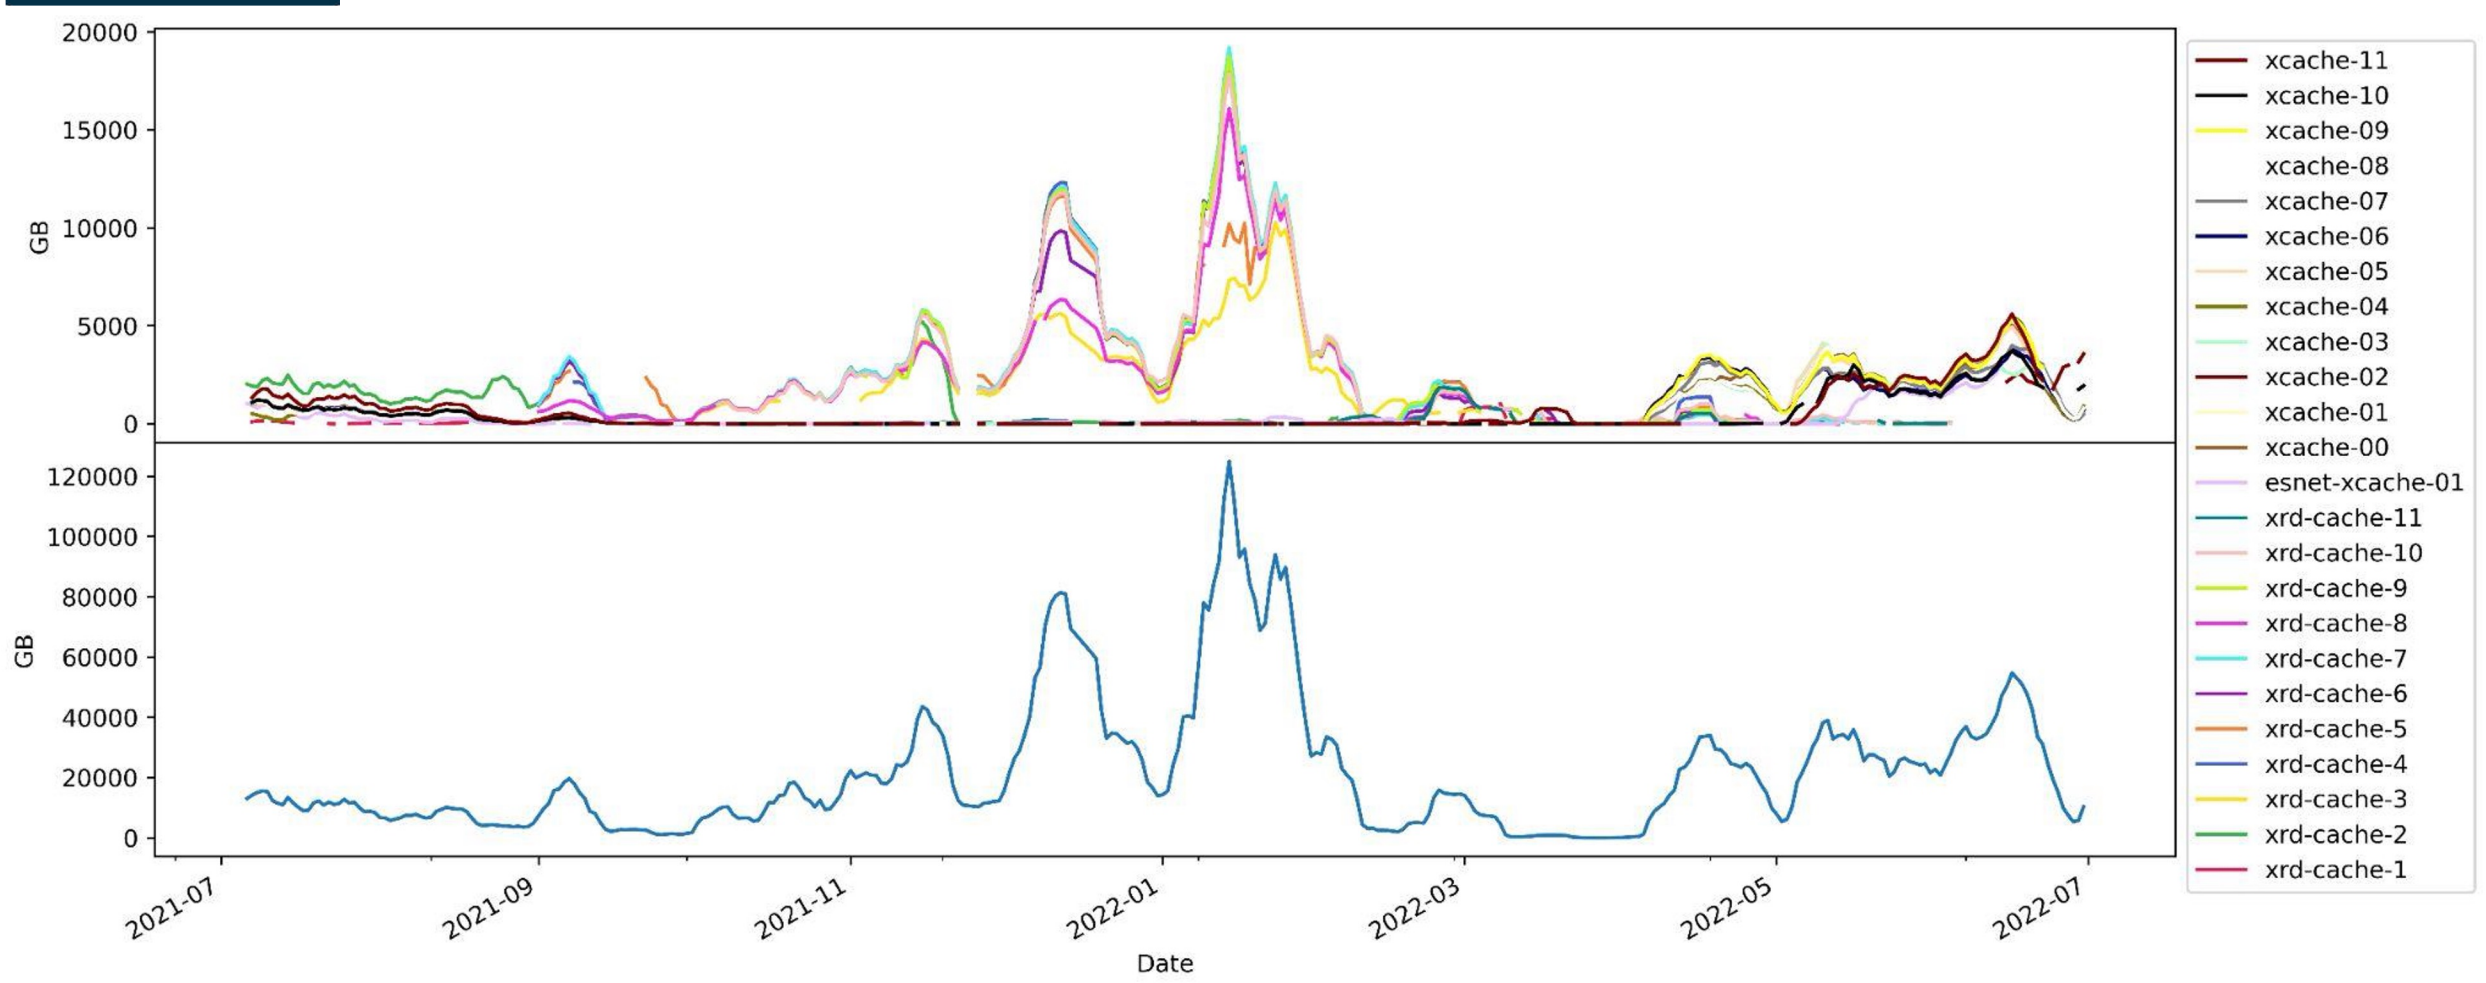

Daily cache miss volume with 10-day moving average

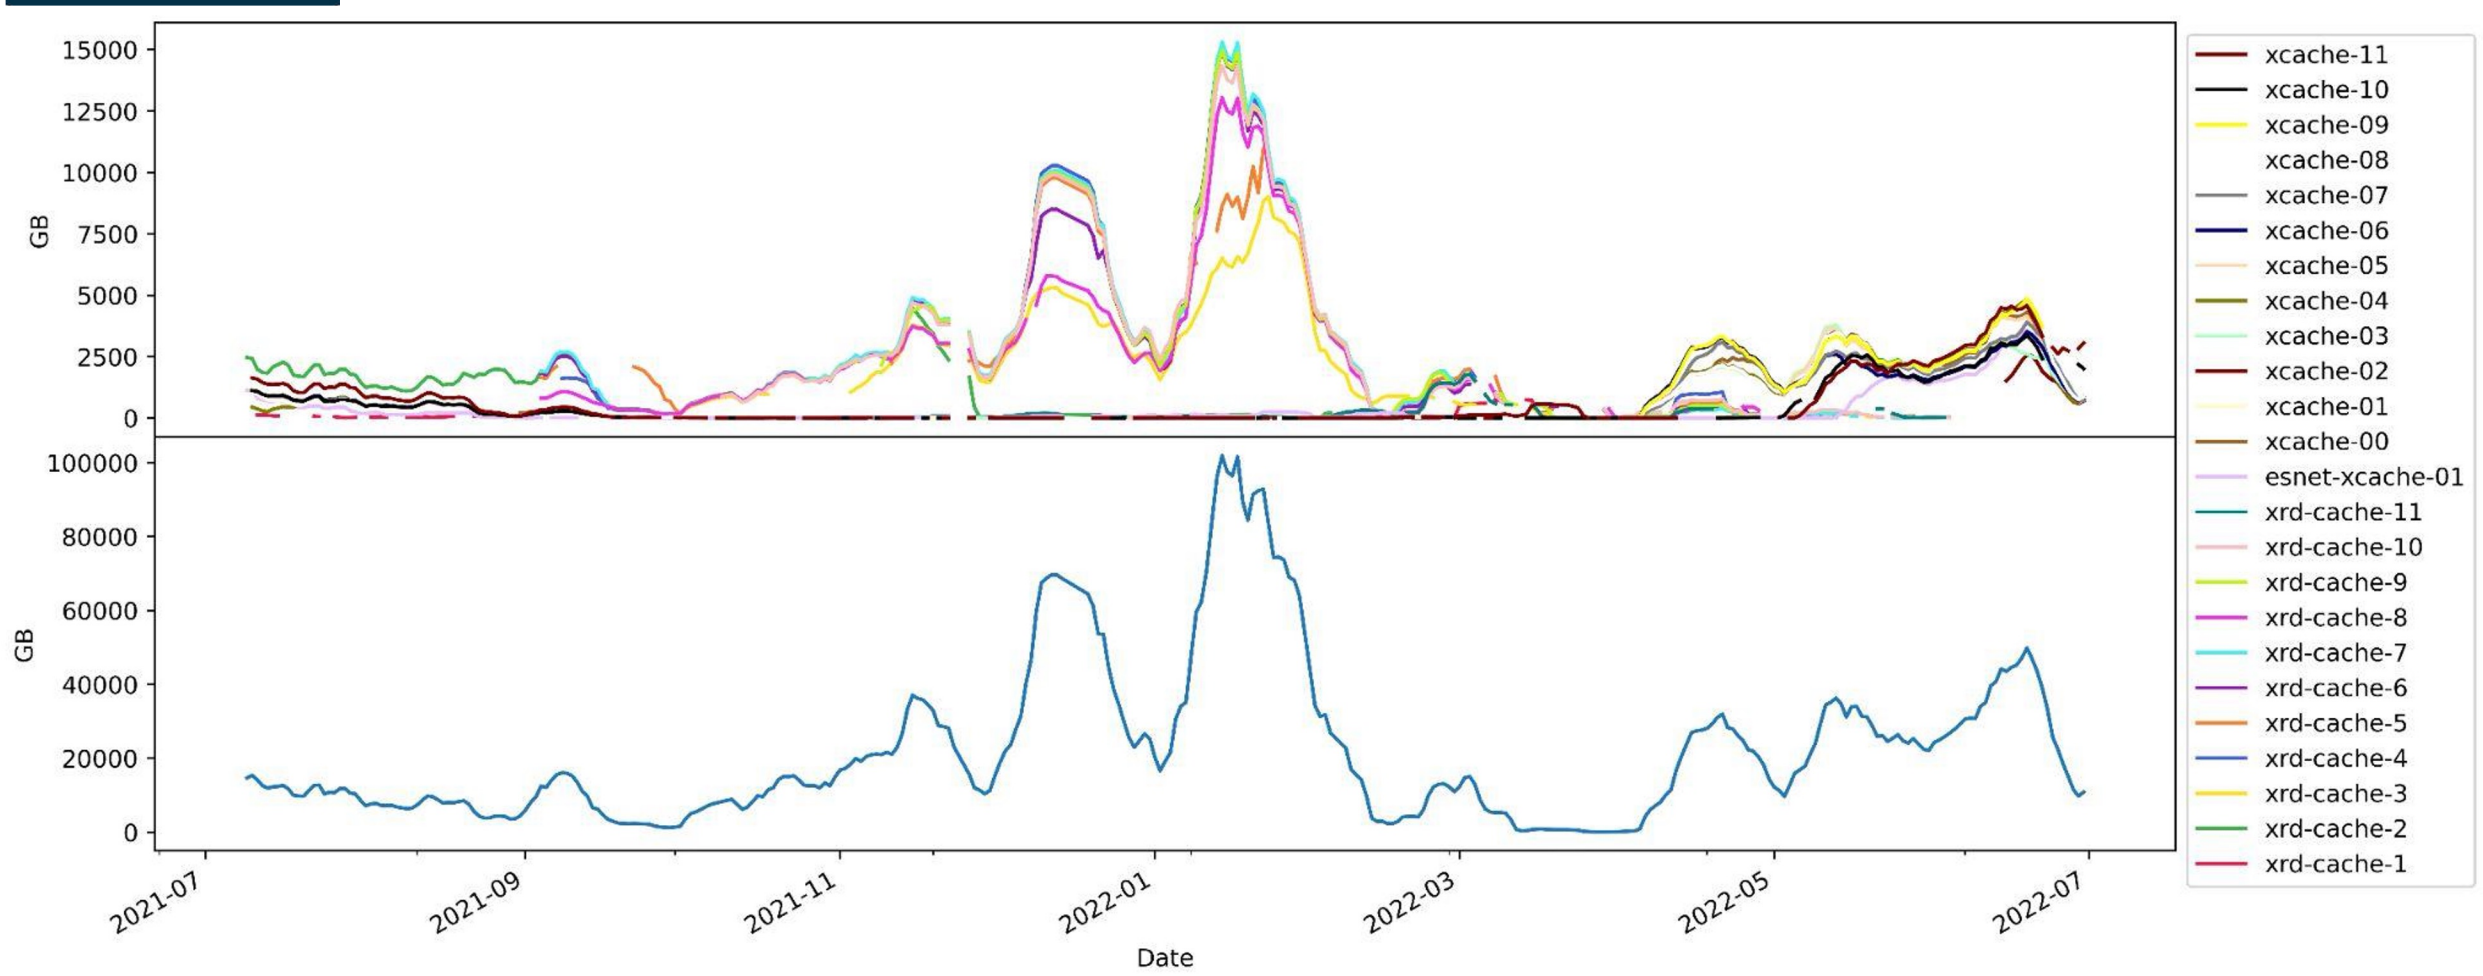

Daily cache miss volume with 14-day moving average

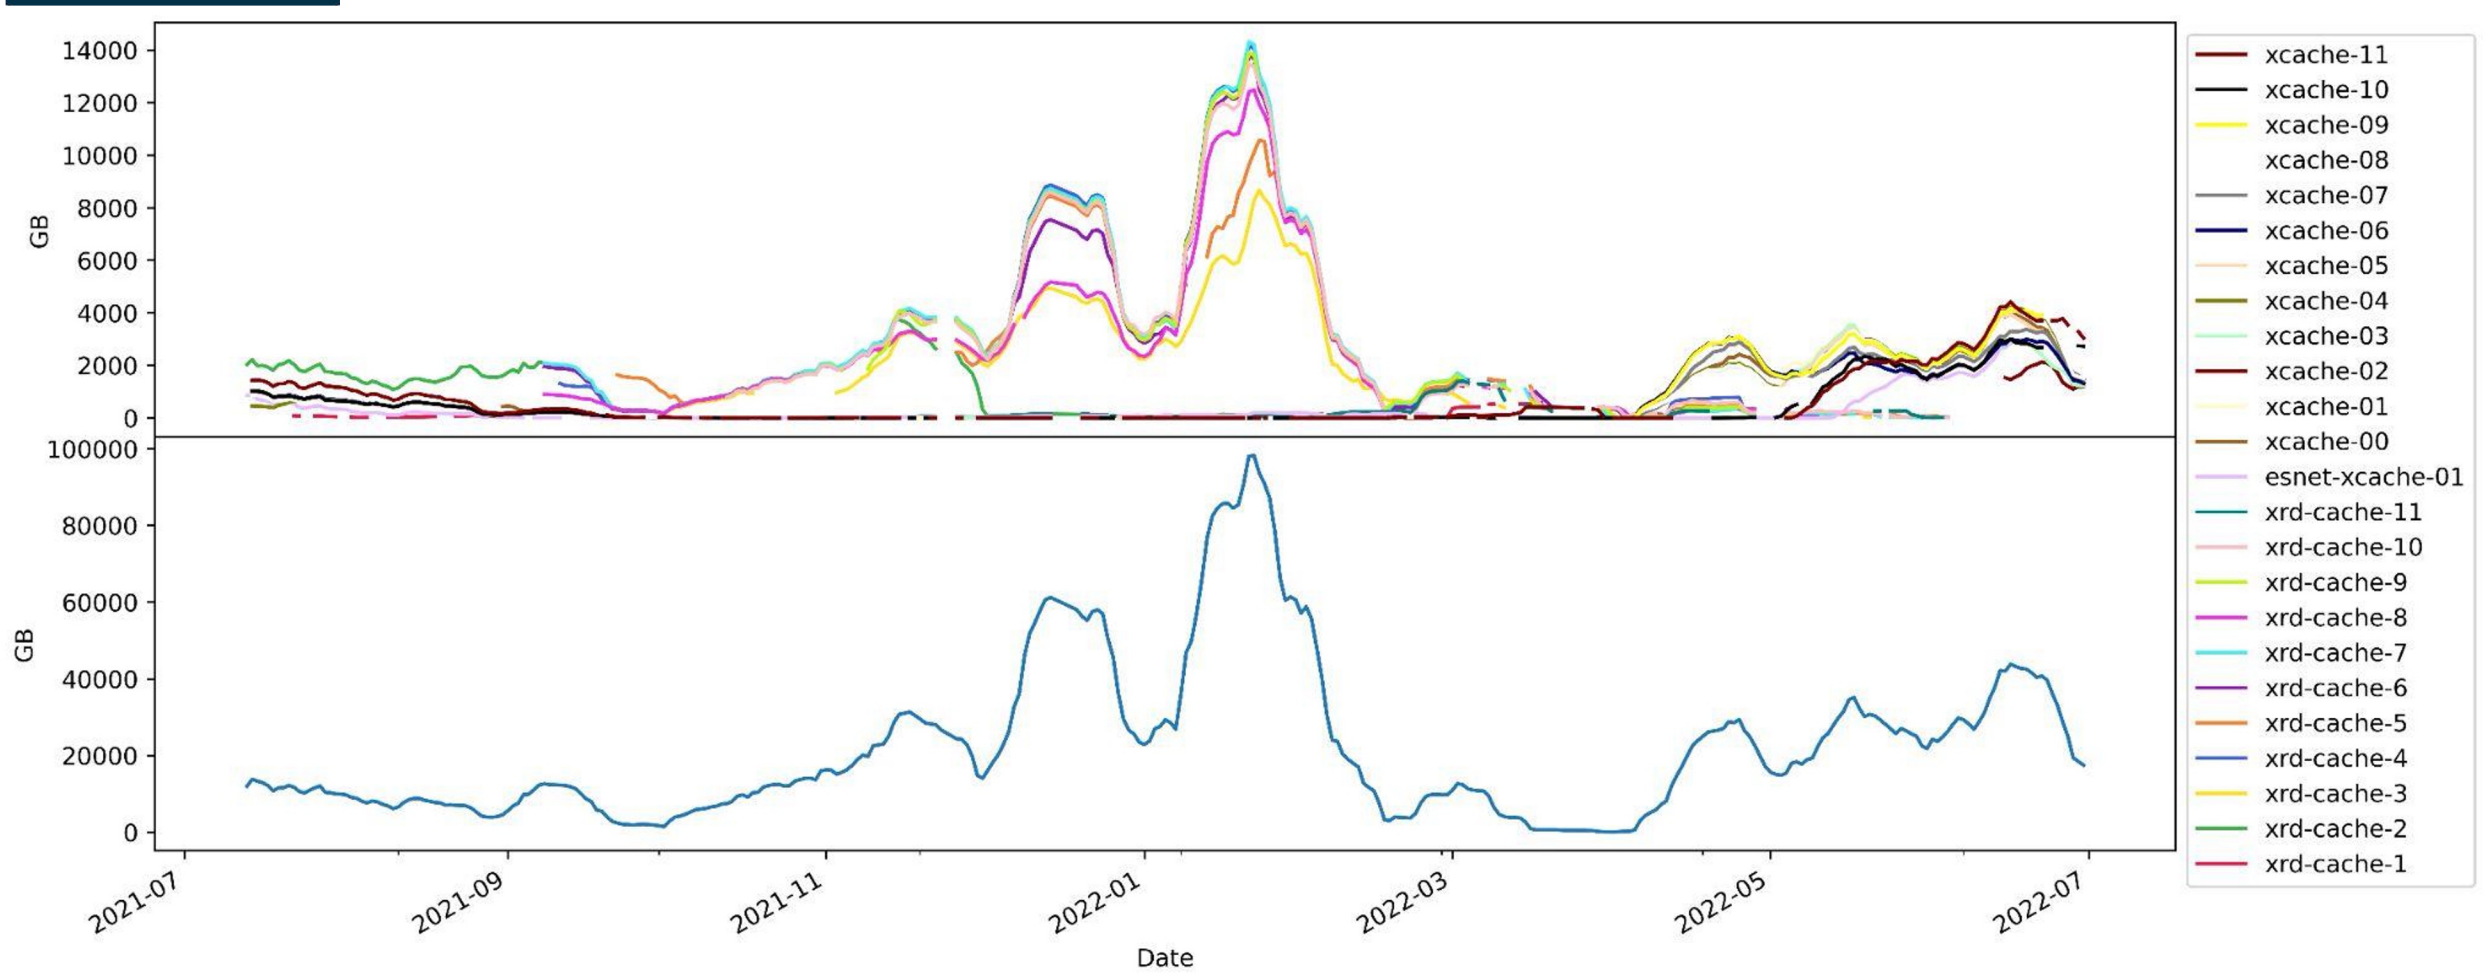

Daily cache miss volume with 21-day moving average

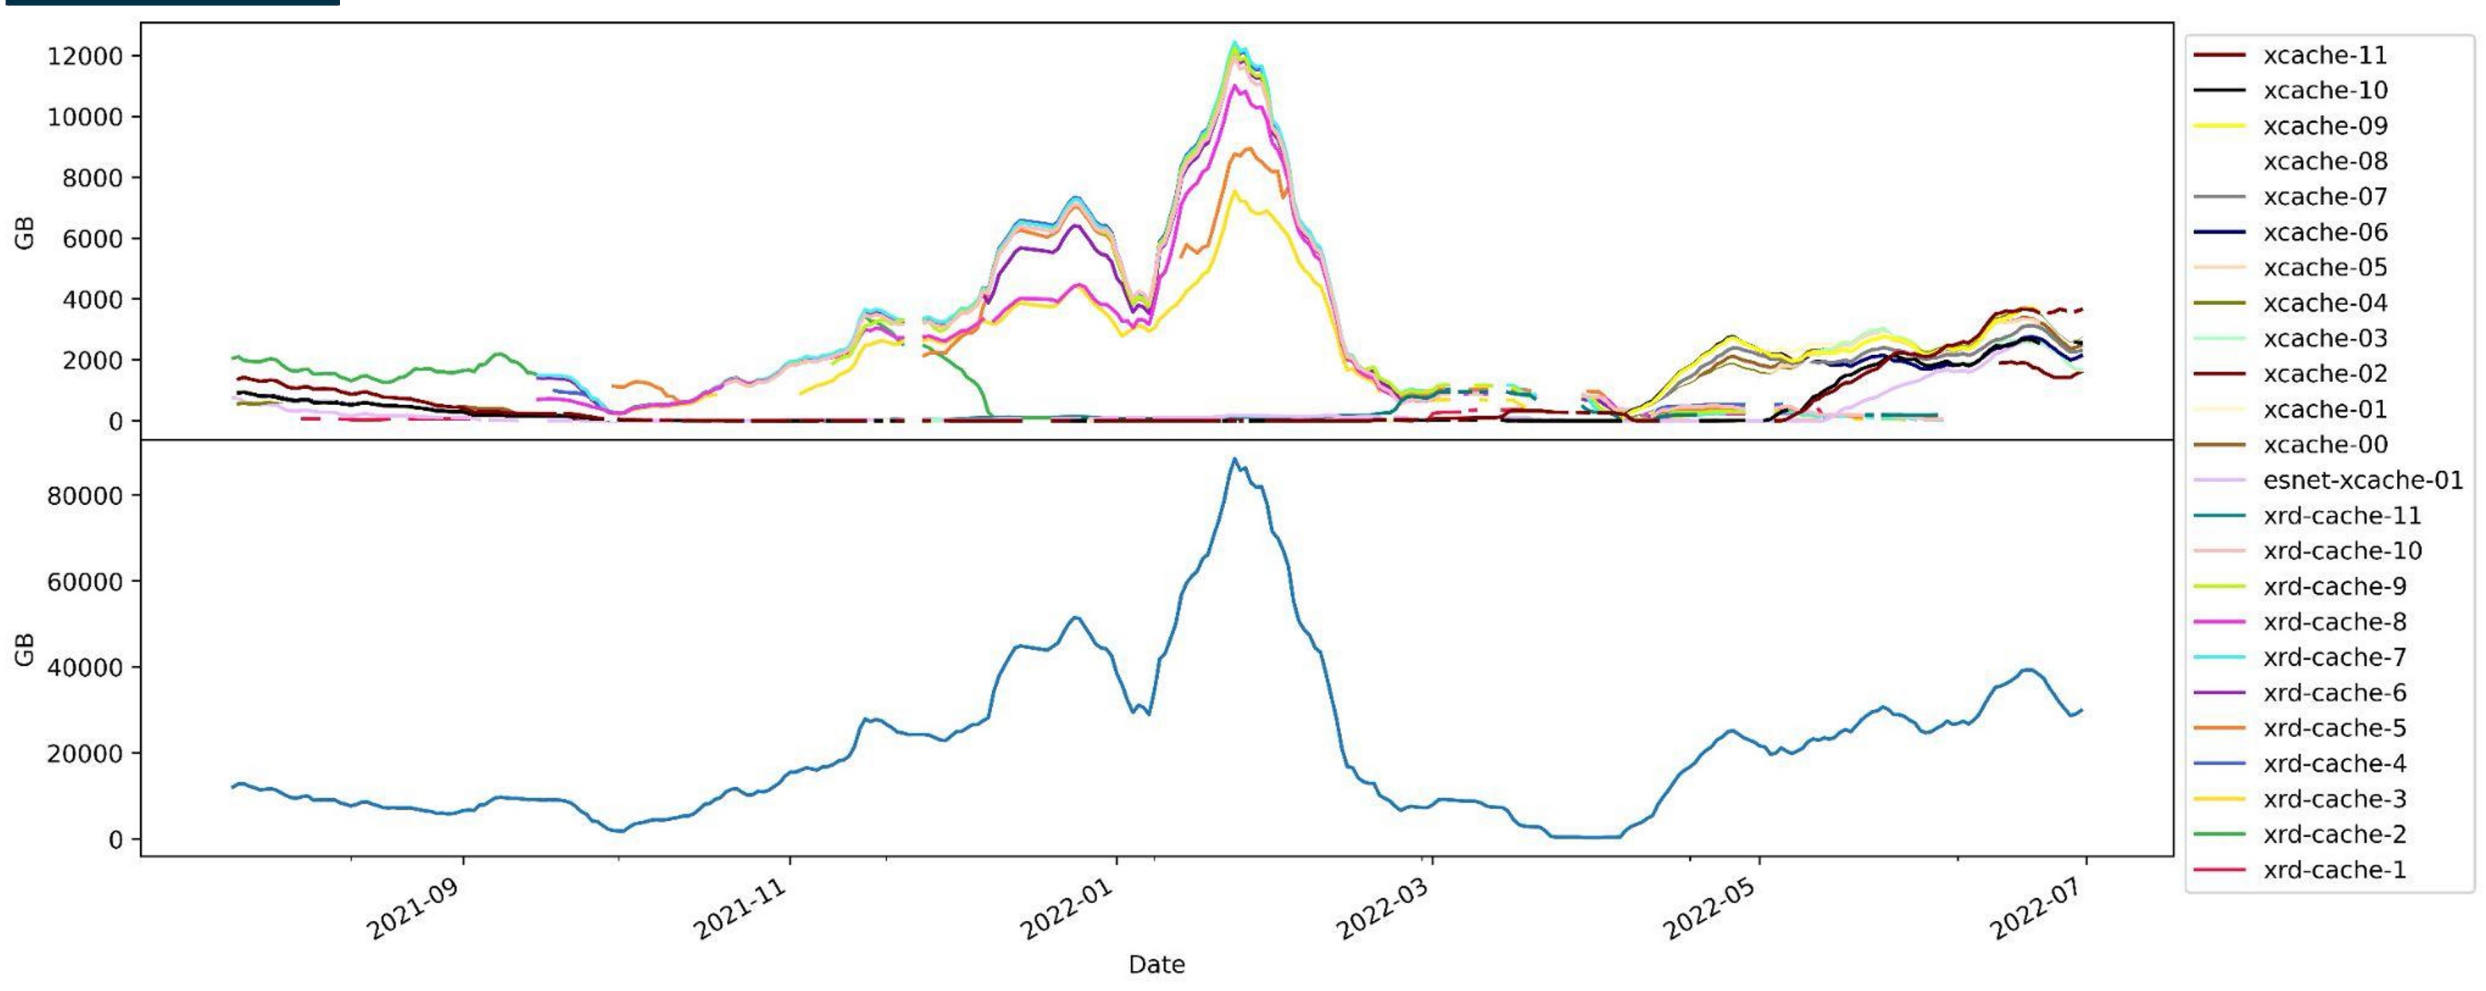

Daily cache miss volume with 28-day moving average

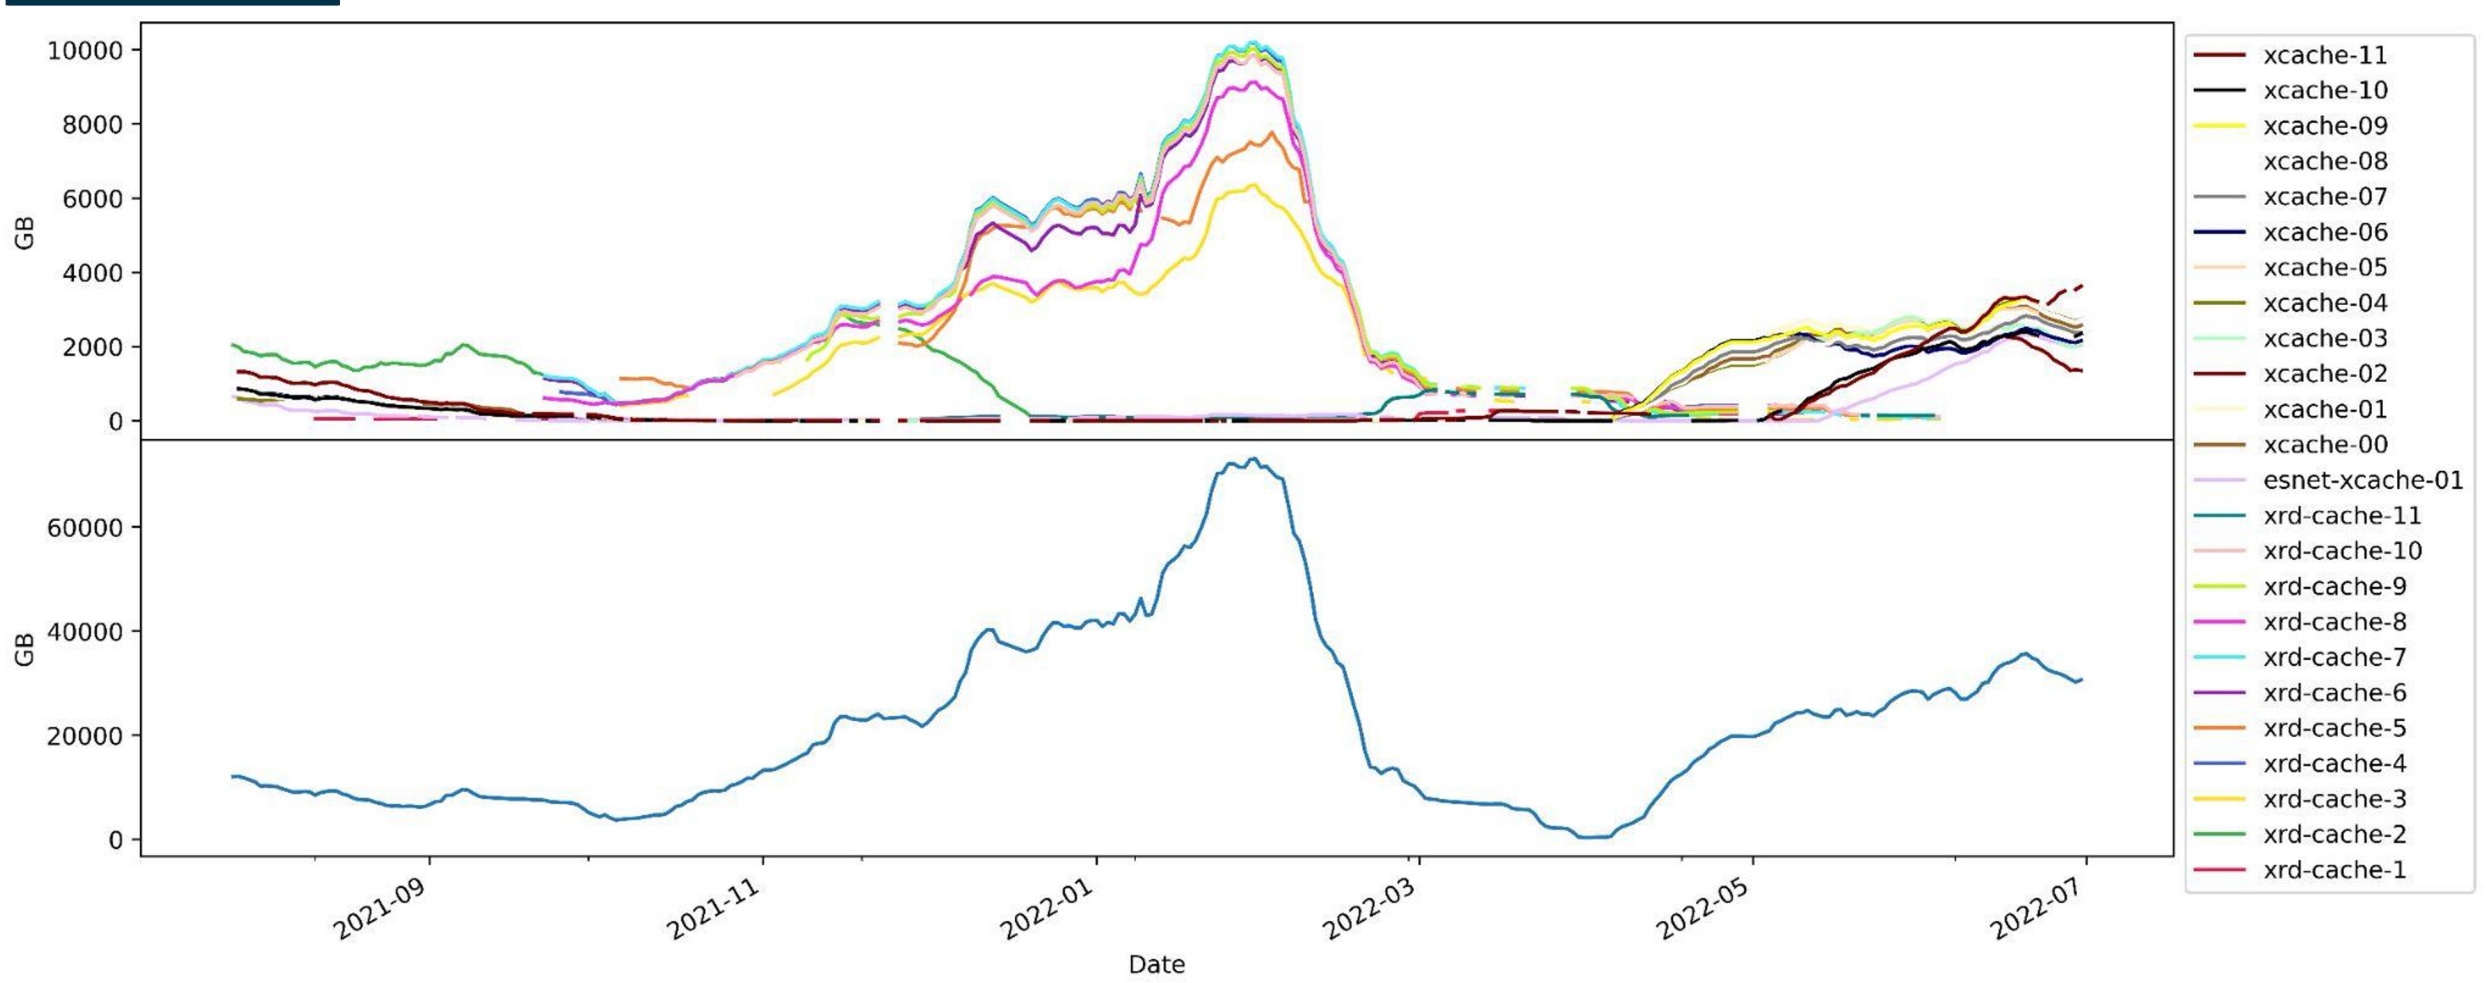

Data re-use and data re-access

- **Data re-use**
  - Successive cache hits on the same data without a cache miss during one day
  - Re-use rates: For the files that are transferred during one day time period, how many times the same data files have been accessed without transfers

- **Data re-access**
  - Cache misses on the same file during one day
  - Xrootd transfers a part of the file each time (not the whole file unless requested)
  - The actual transferred data blocks may be different (or the same) for each transfer
  - Data re-access indicates that the same file has been requested, “transferred”, and accessed
  - Re-access rates: How many times the same file (name) has been transferred during one day time period.

Daily data re-use rates (in log scale)

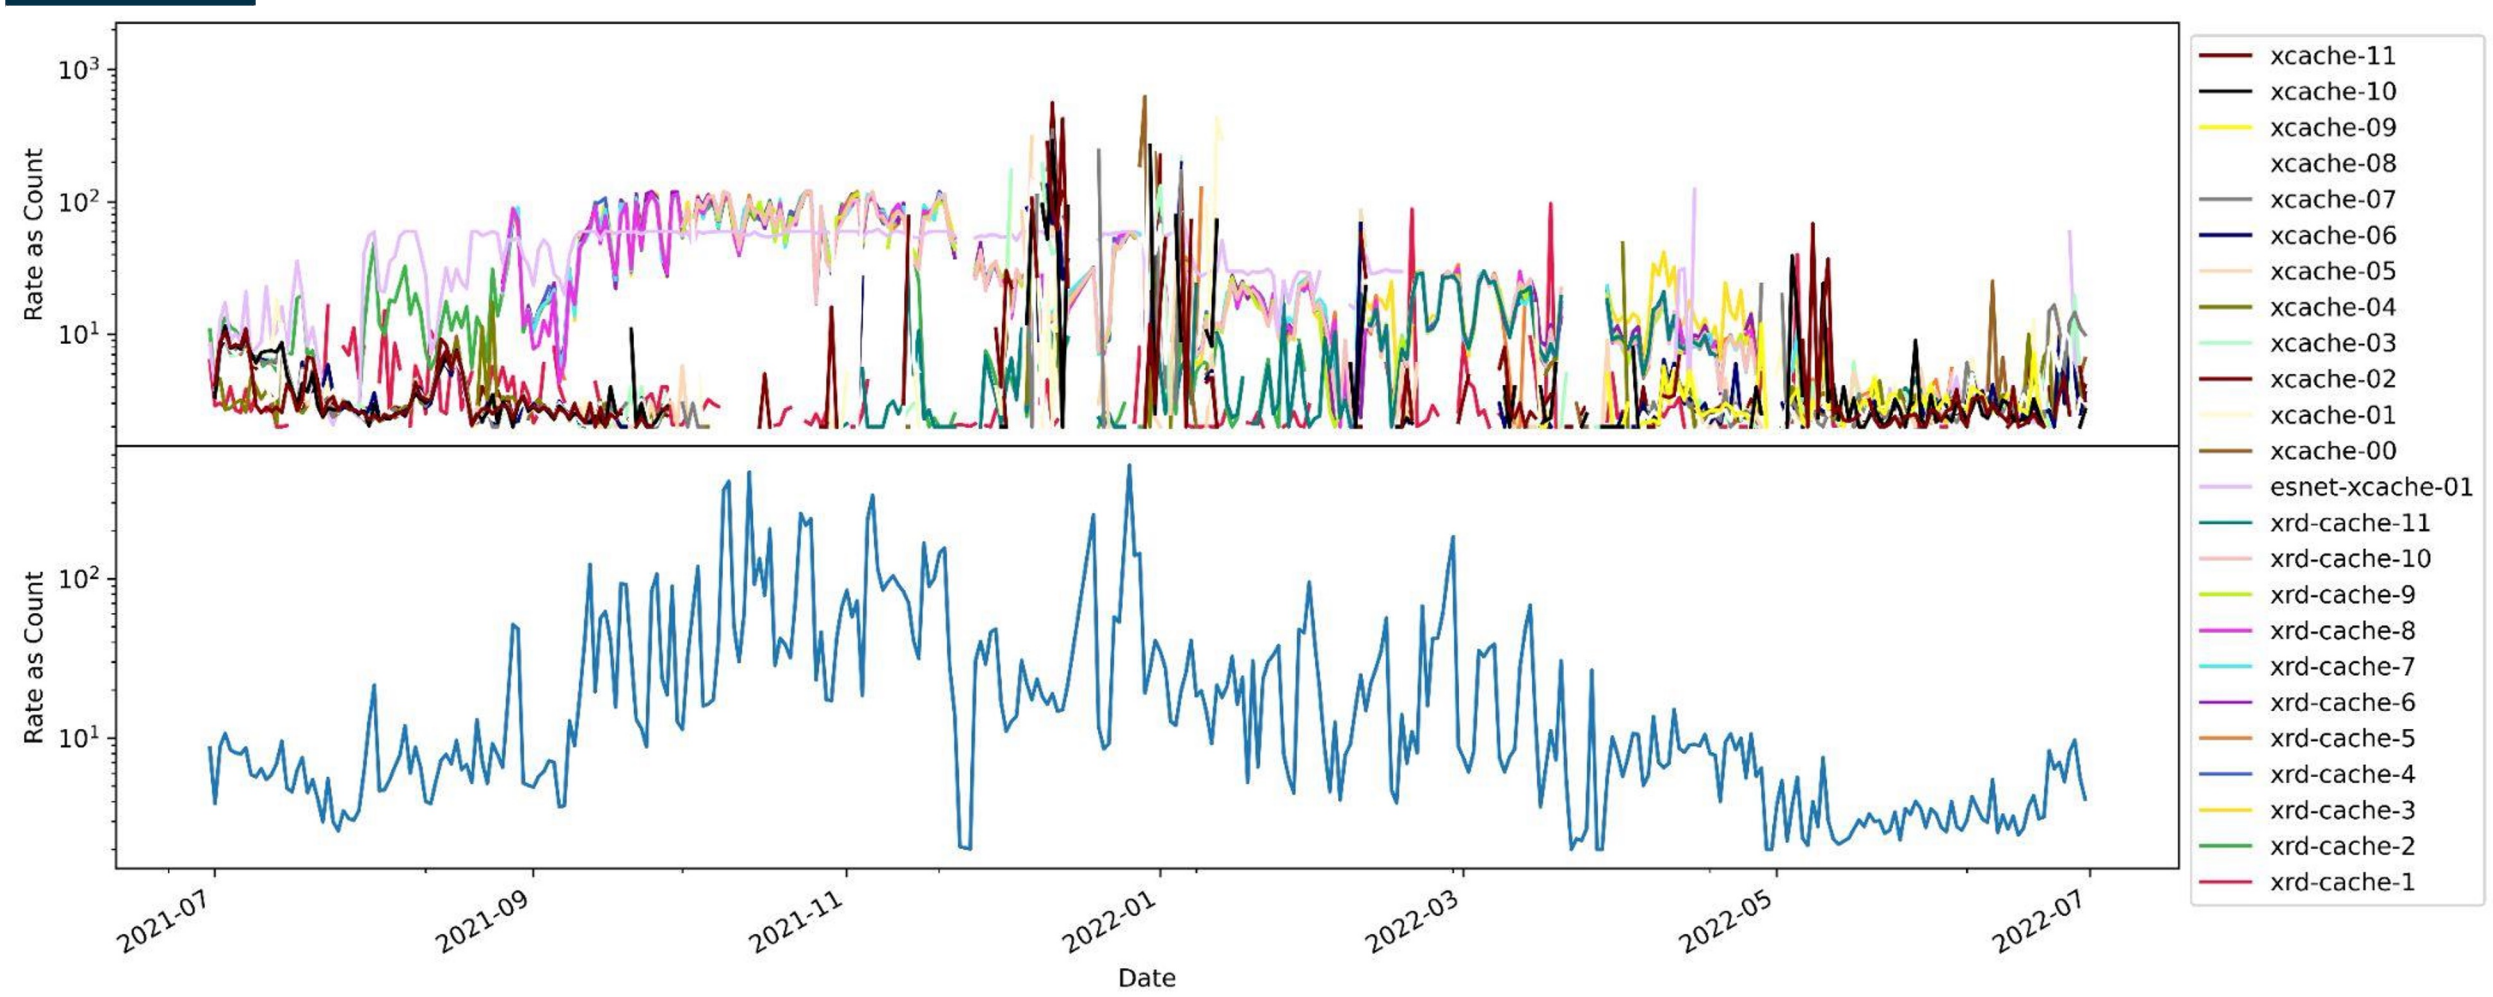

Daily data re-use rates with 2-day moving average (log scale)

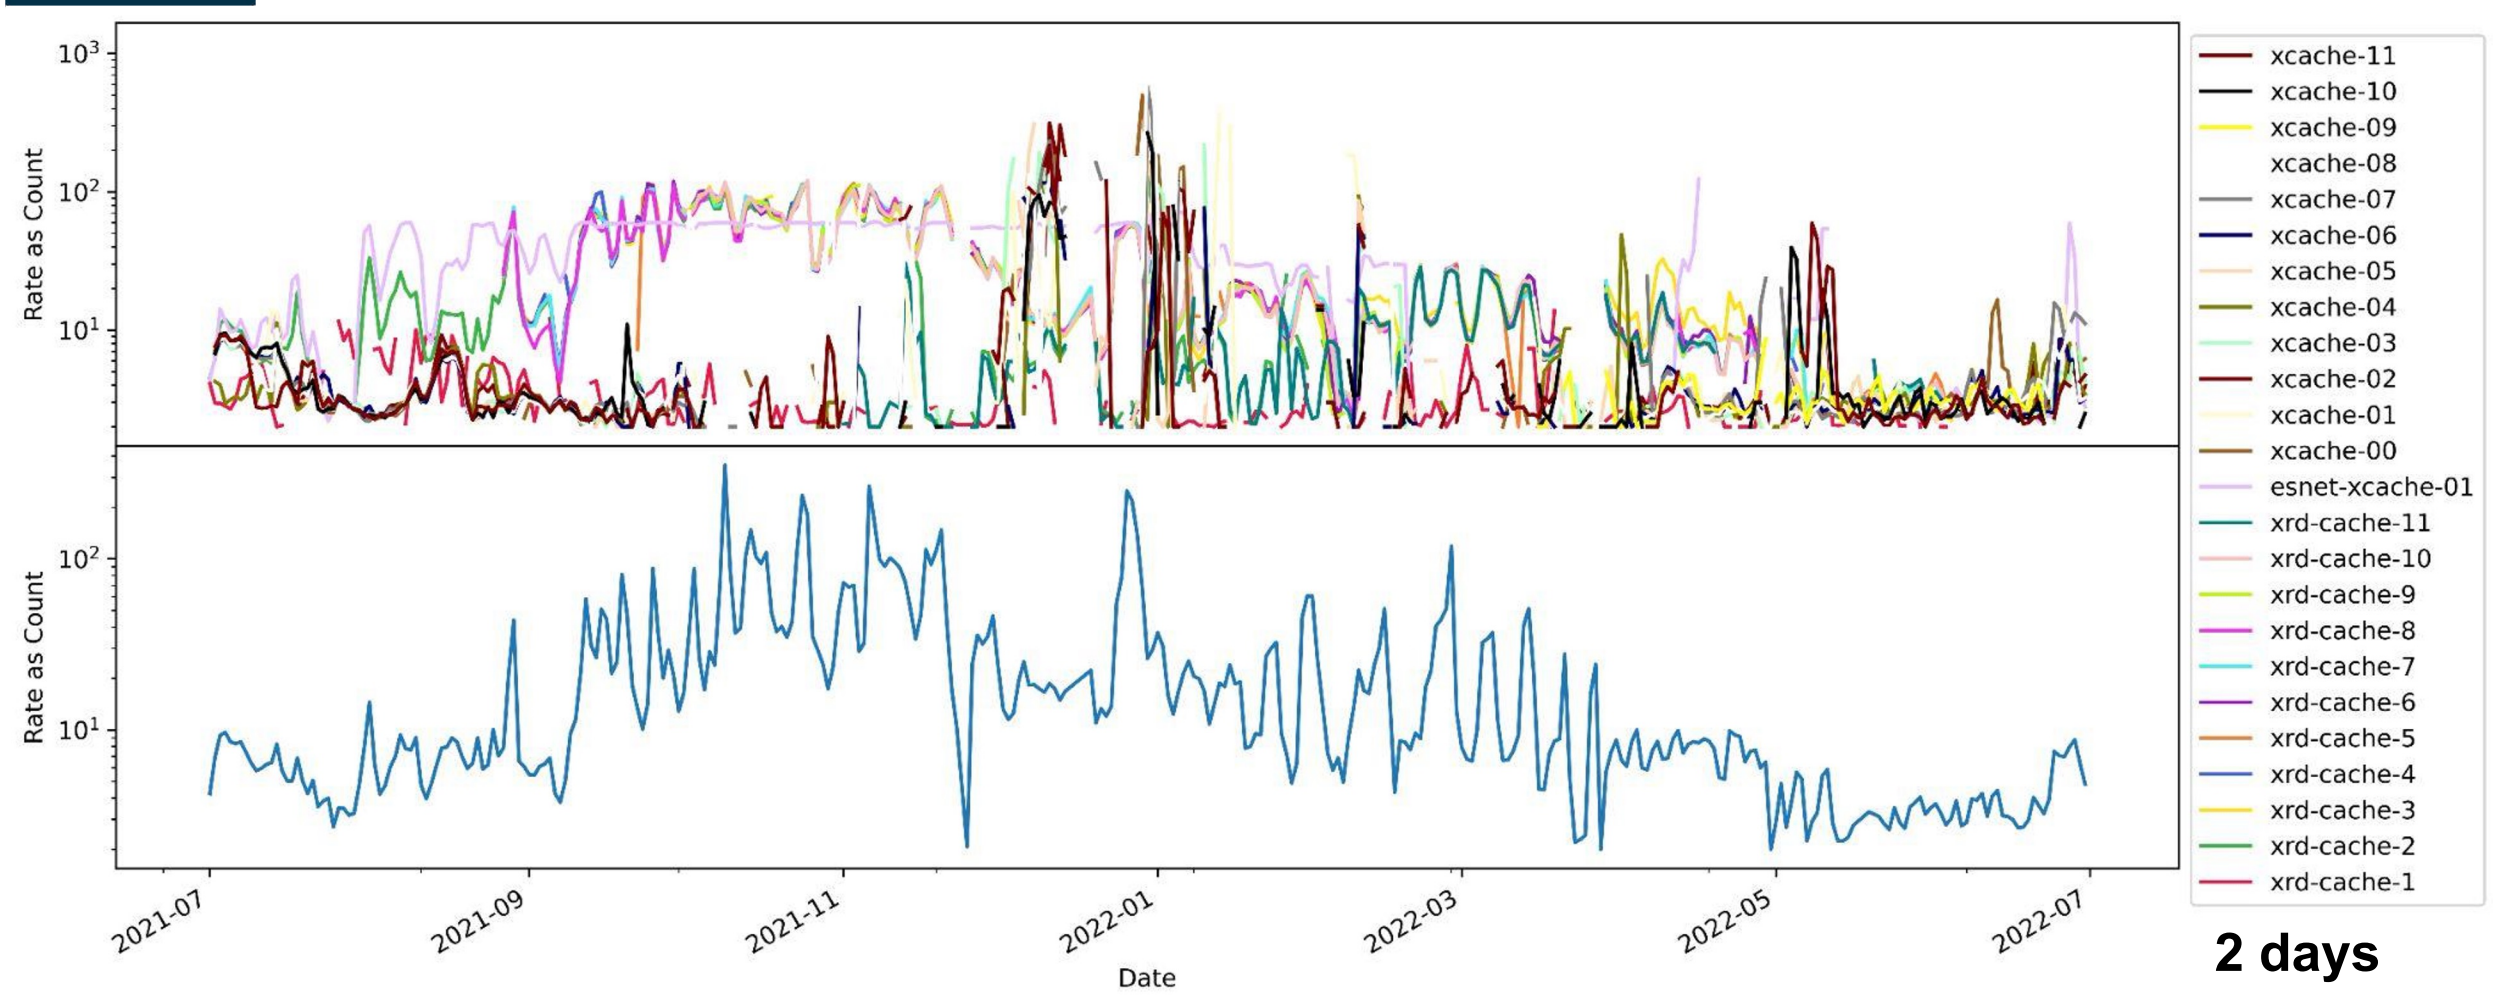

2 days

Daily data re-use rates with 3-day moving average (log scale)

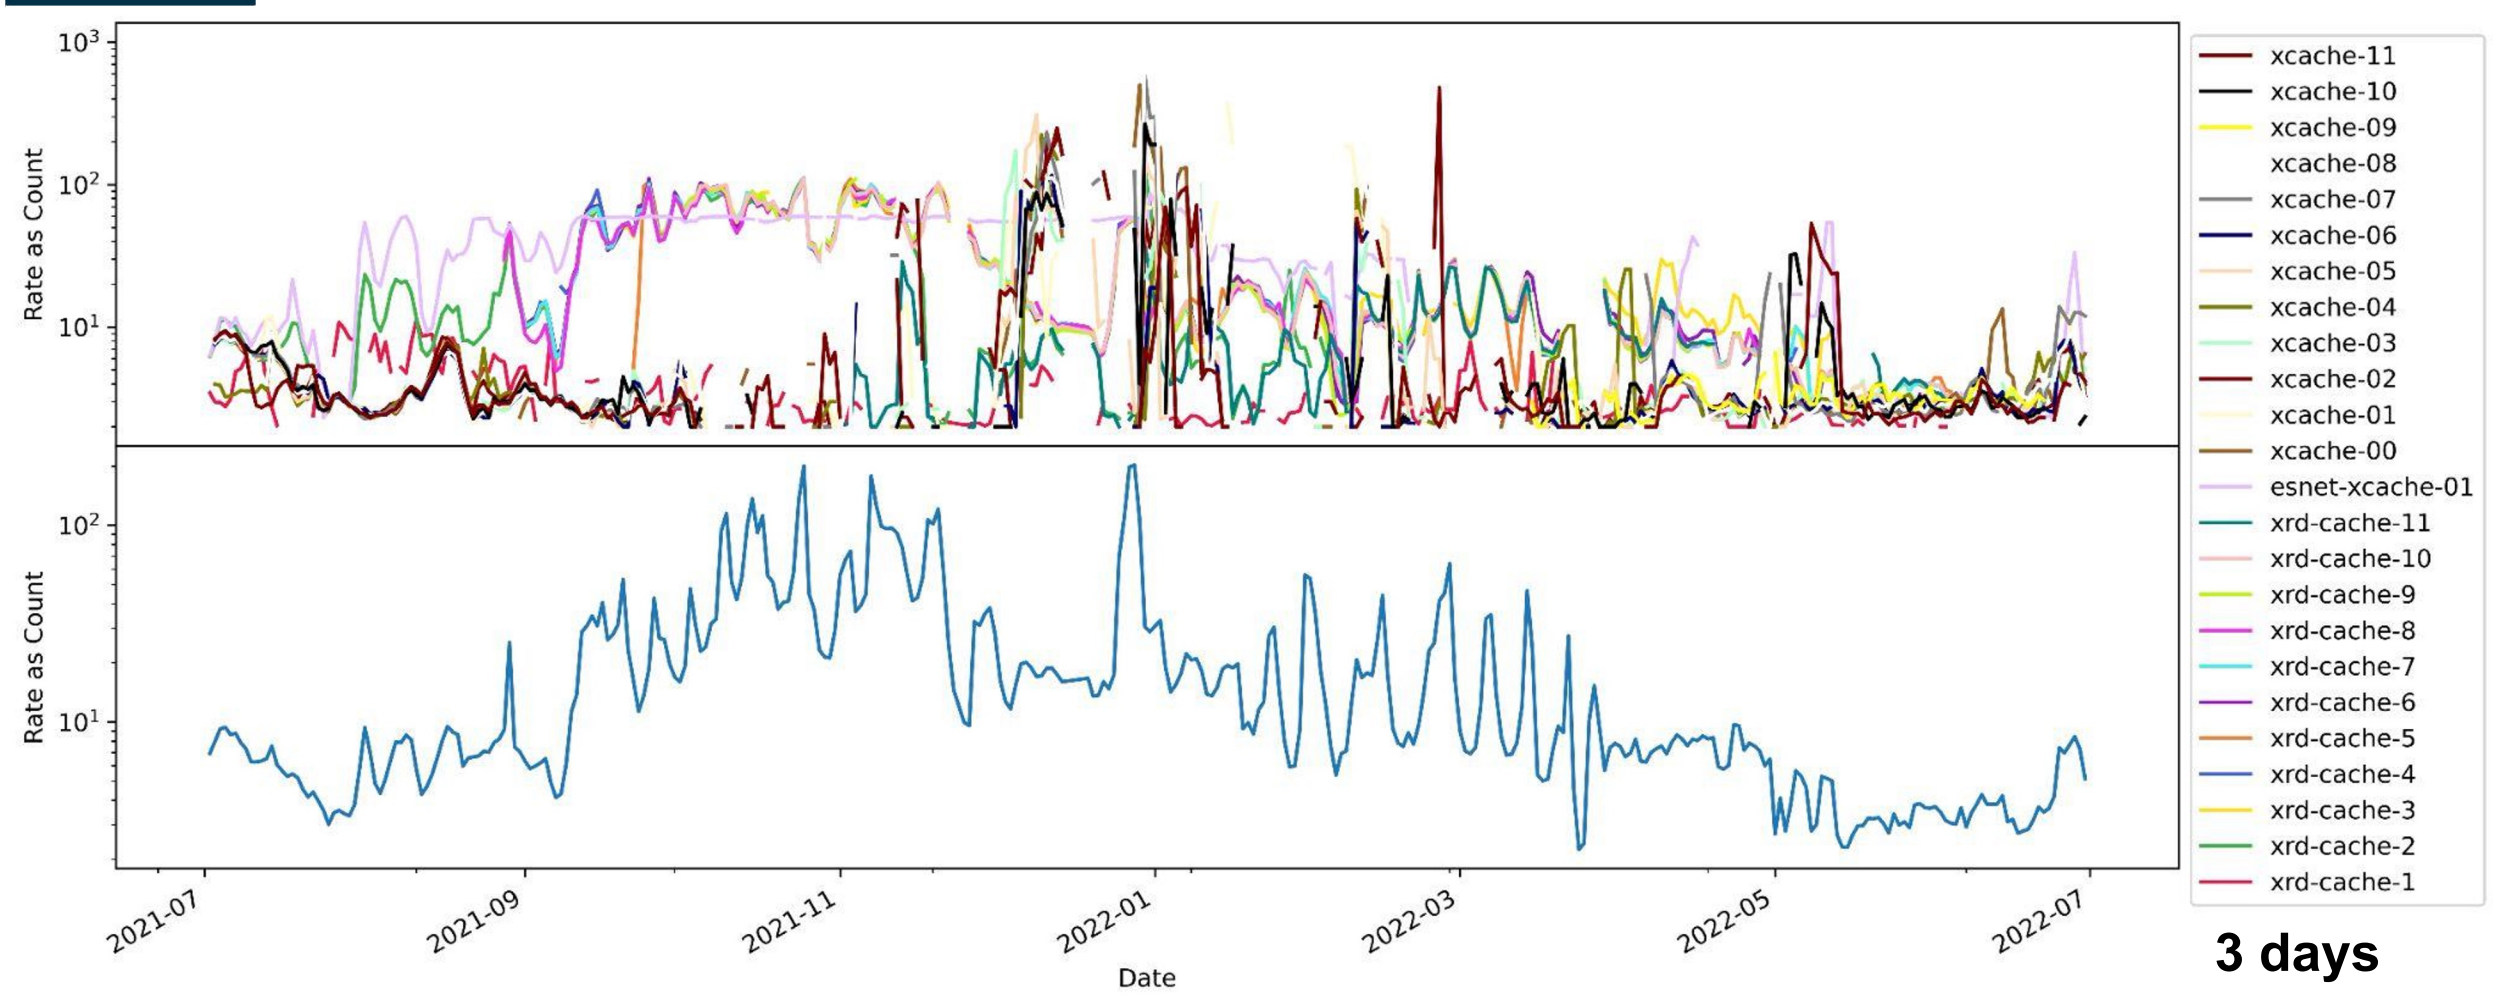

3 days

Daily data re-use rates with 5-day moving average (log scale)

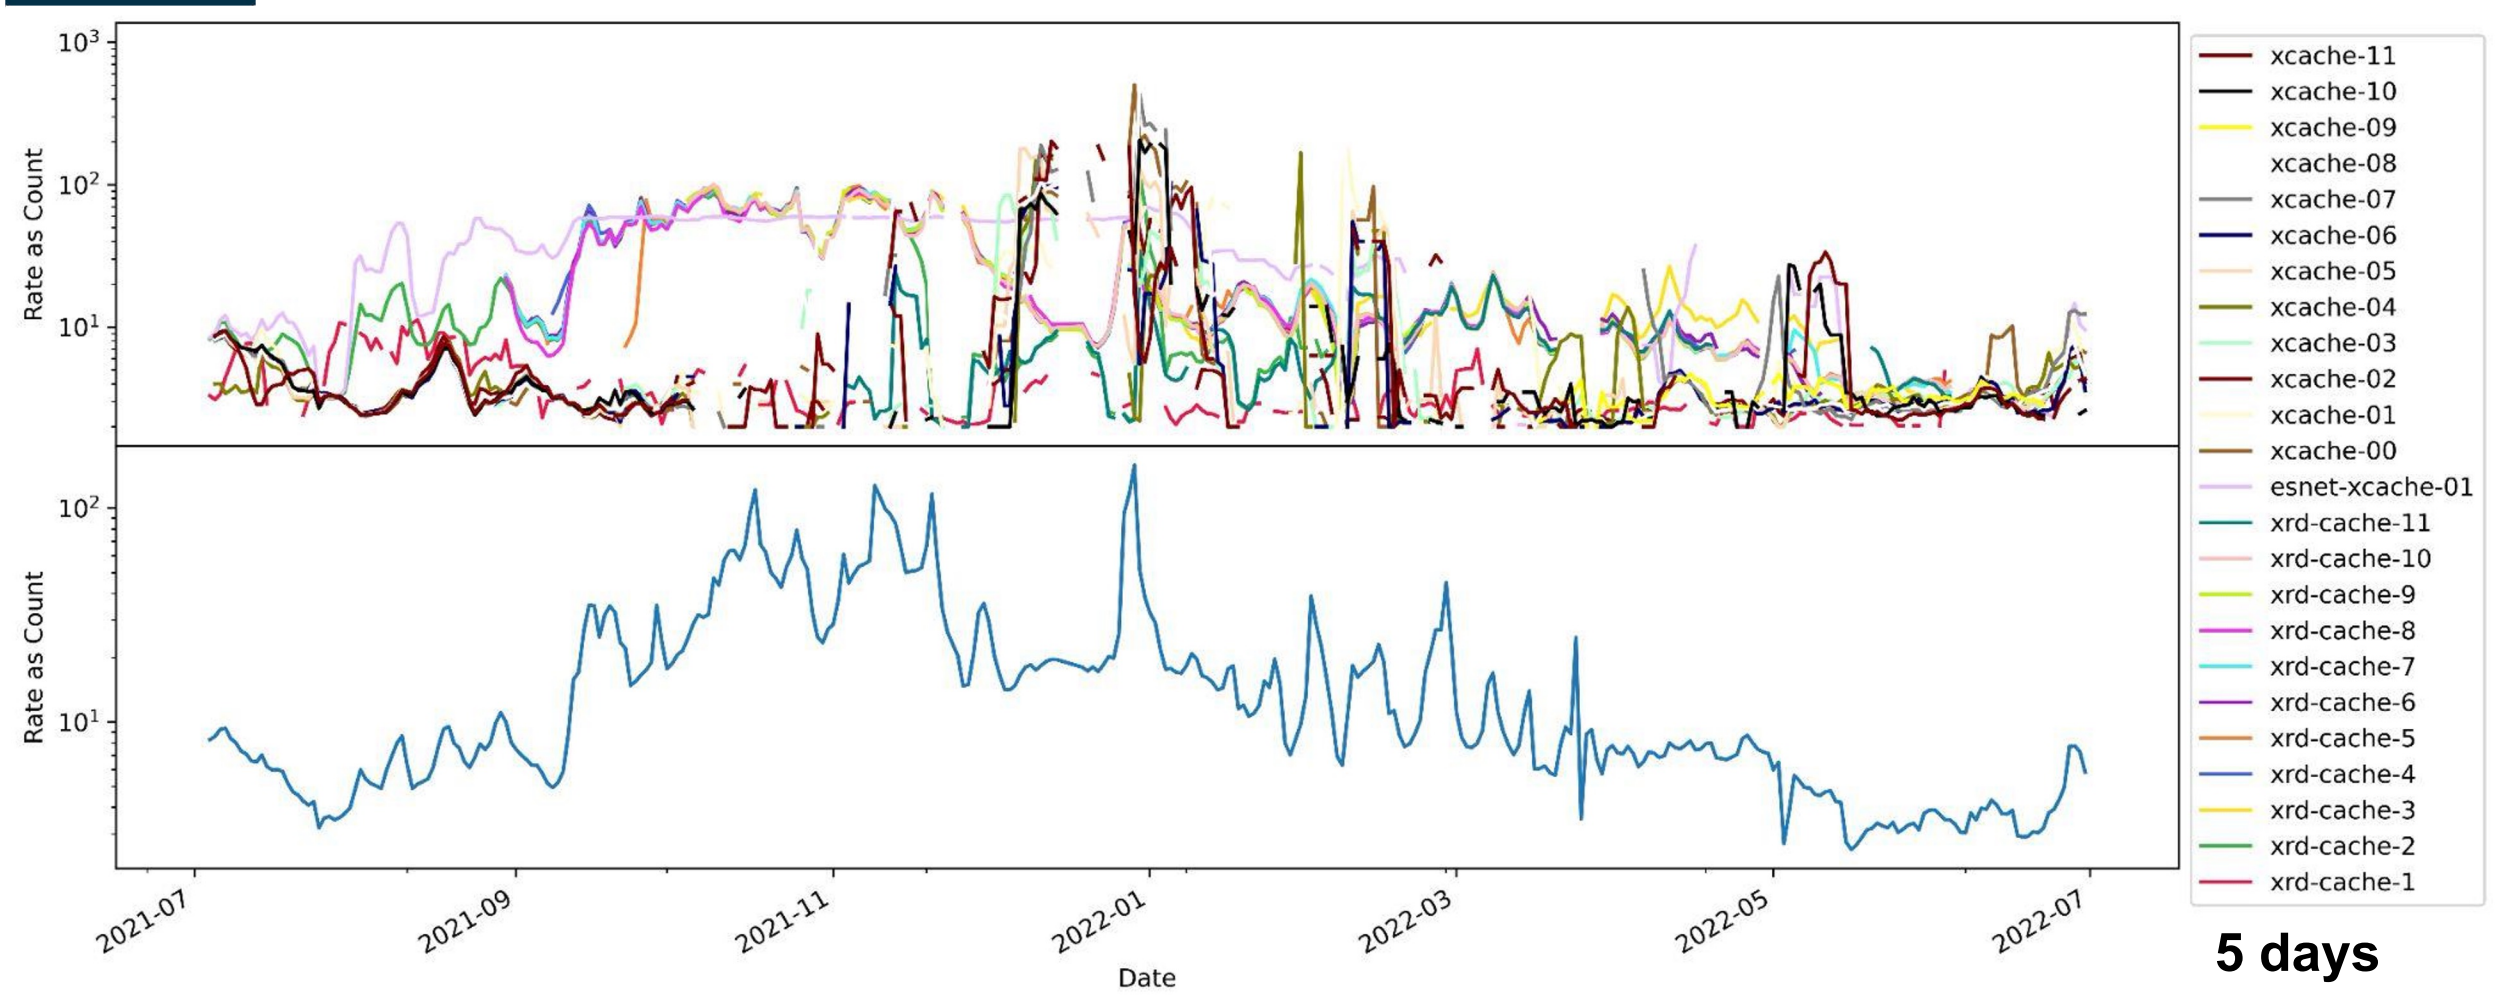

5 days

Daily data re-use rates with 7-day moving average (log scale)

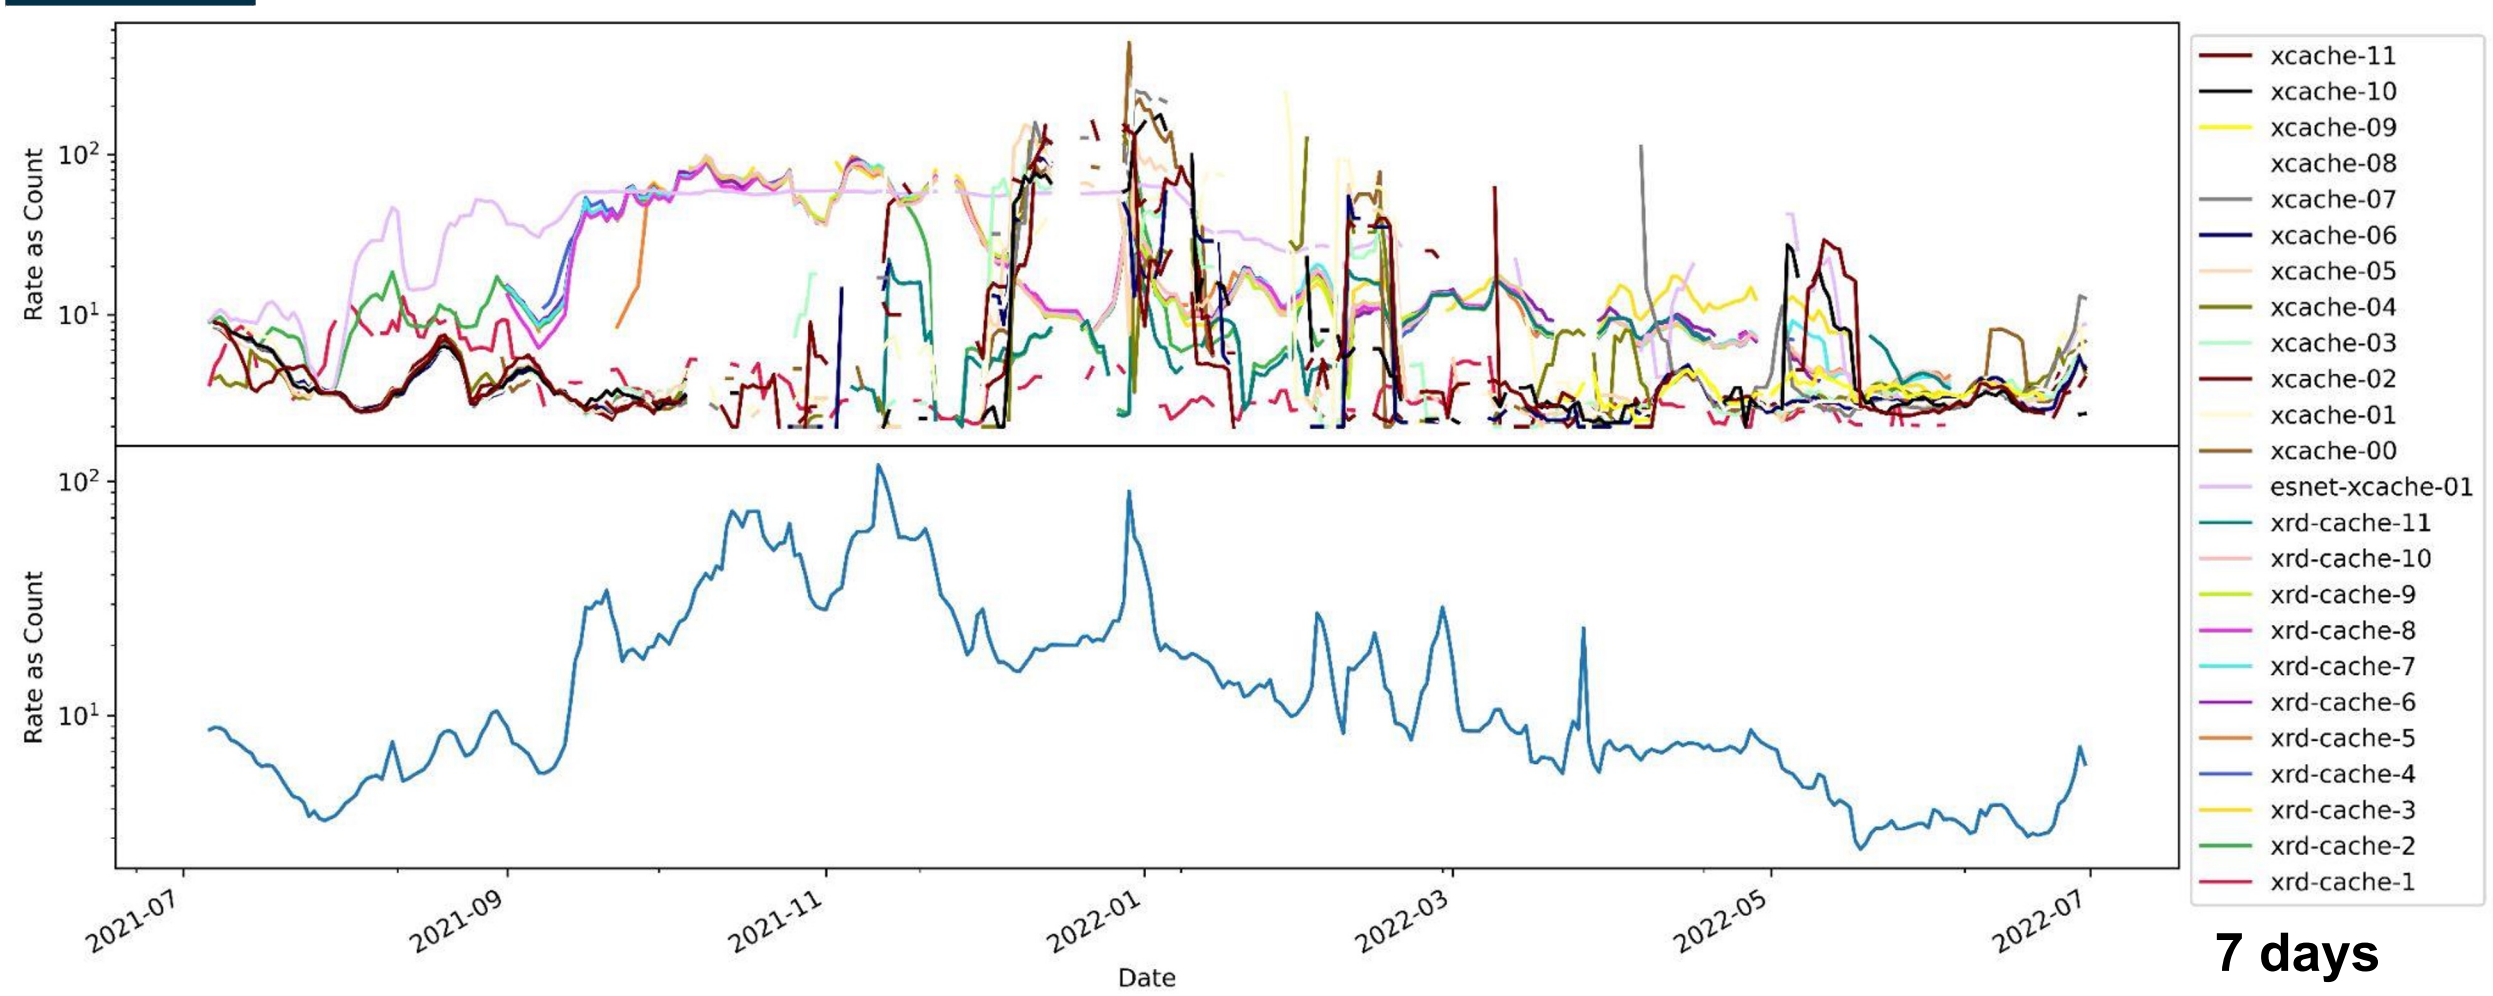

7 days

Daily data re-use rates with 10-day moving average (log scale)

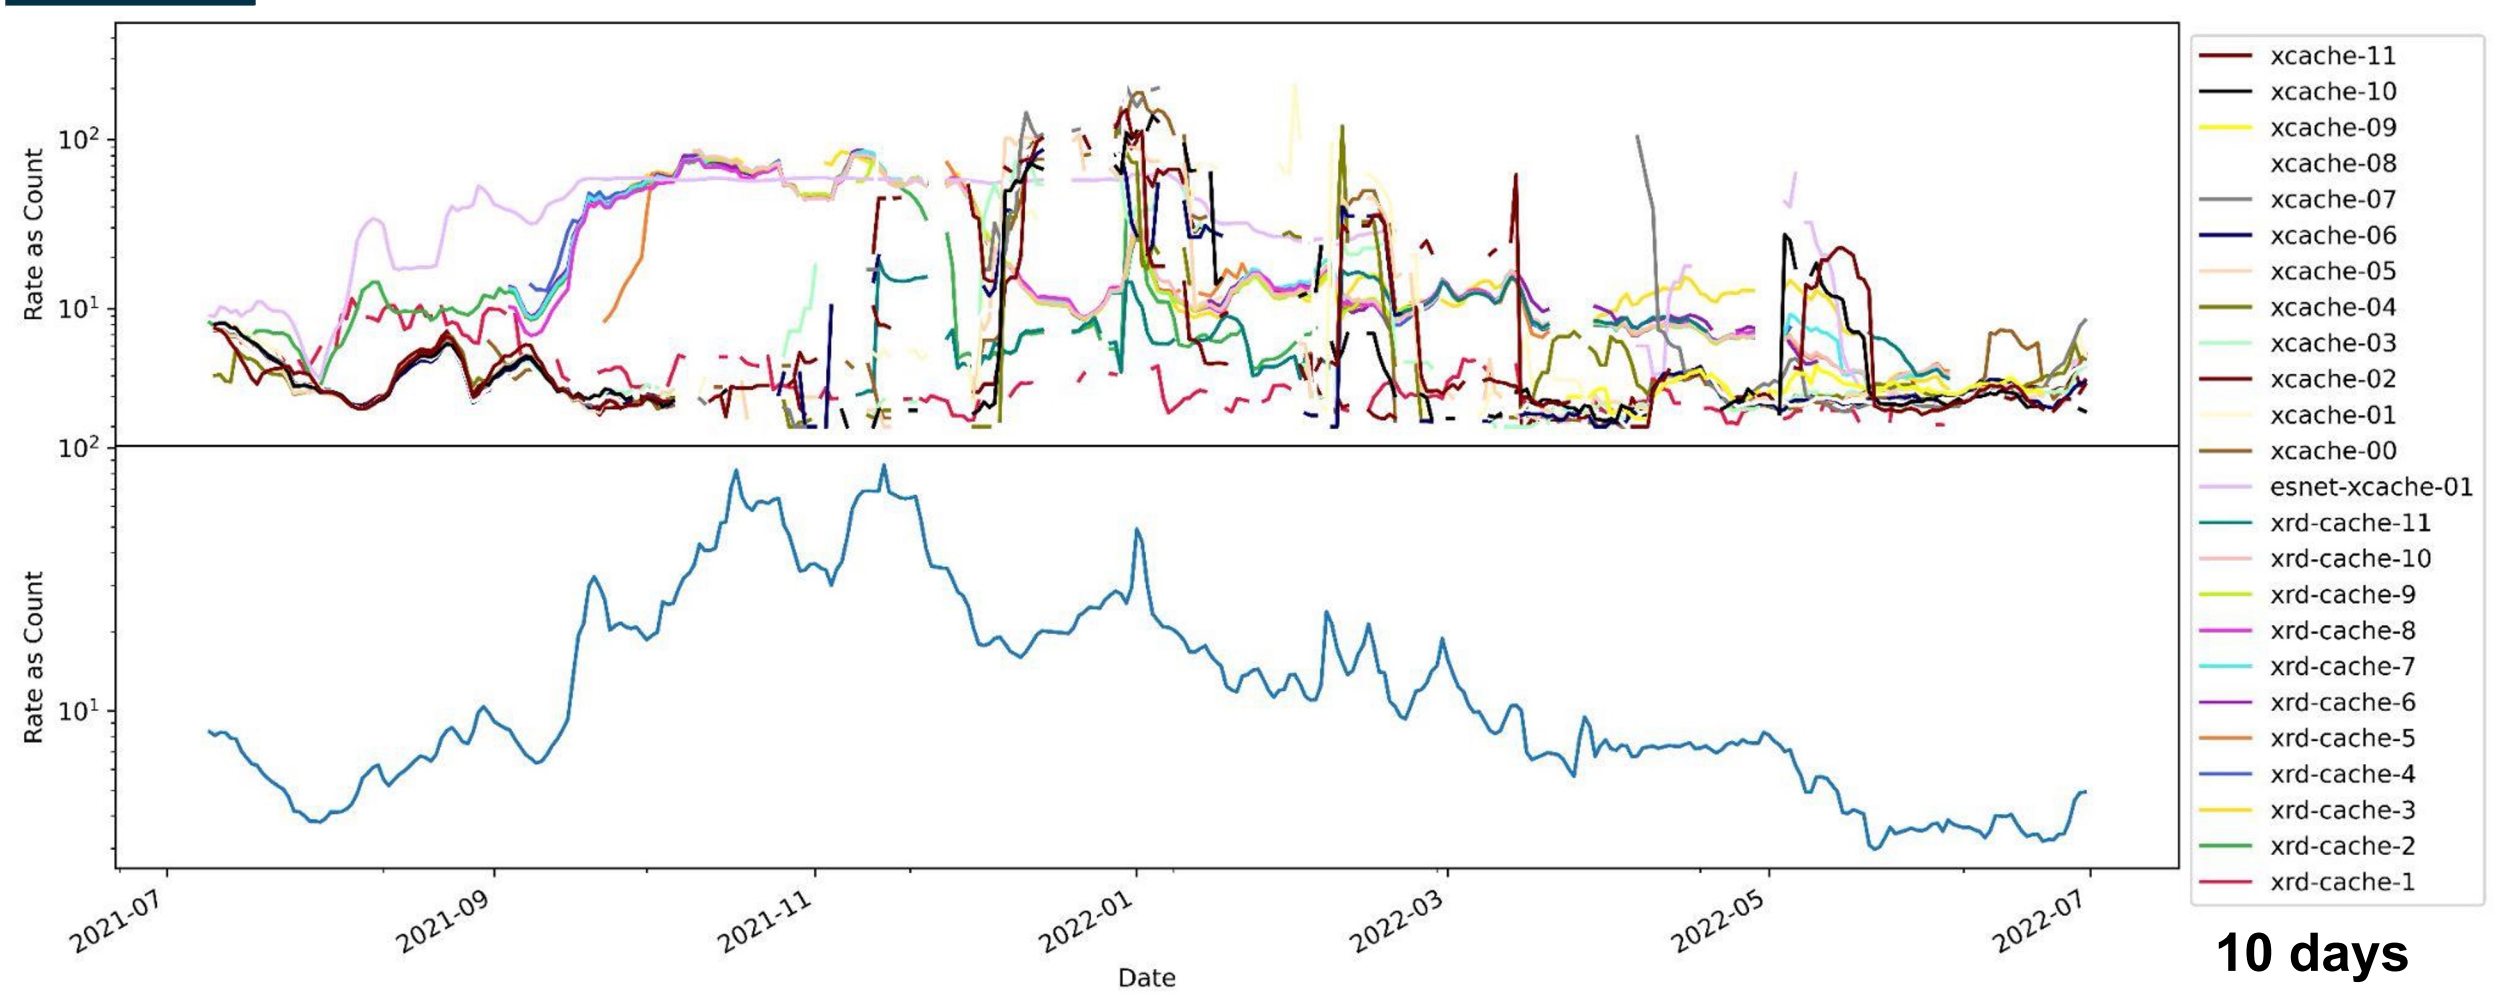

10 days

Daily data re-use rates with 14-day moving average (log scale)

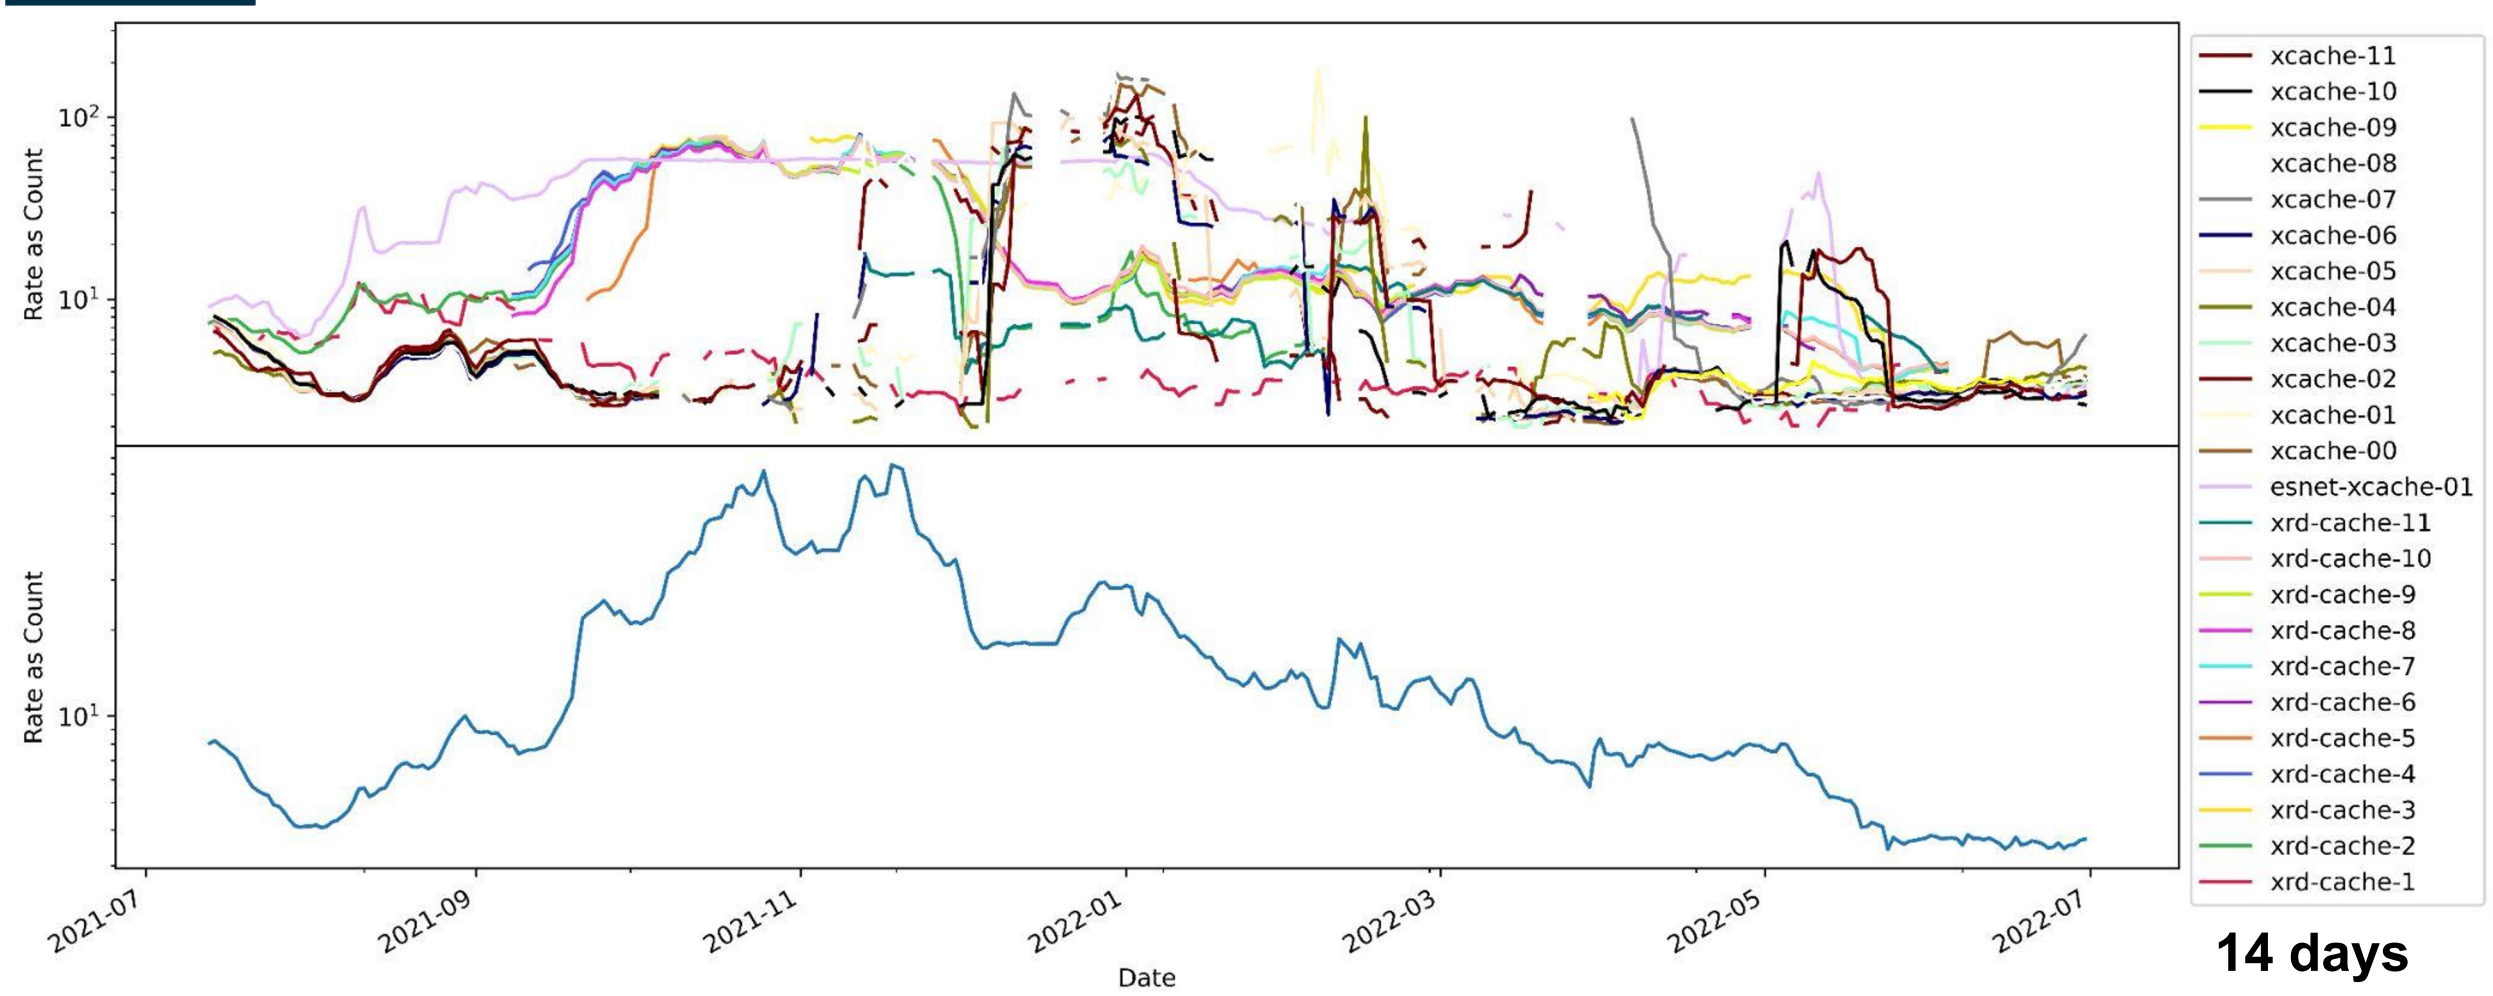

14 days

Daily data re-use rates with 21-day moving average (log scale)

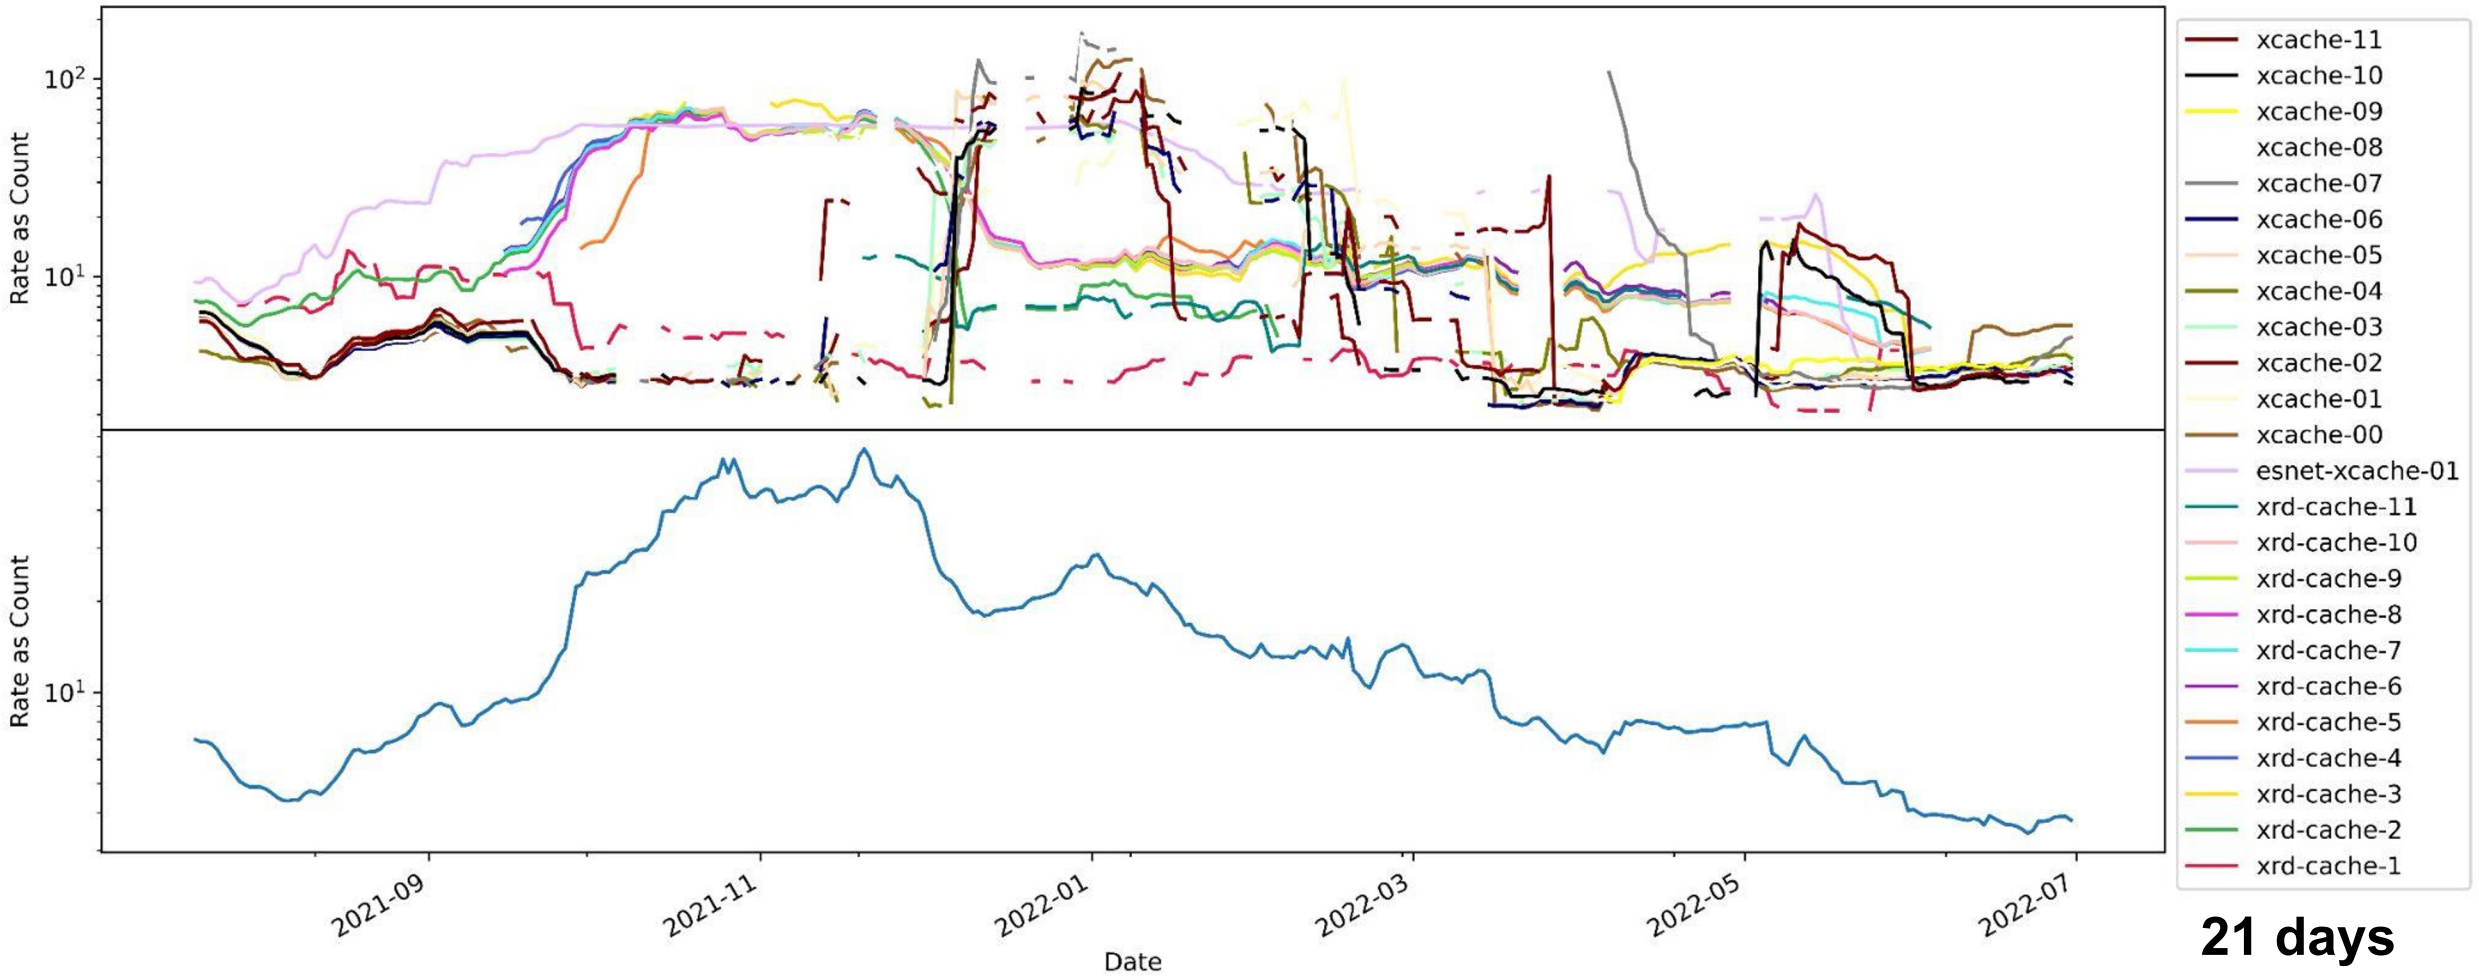

21 days

Daily data re-use rates with 28-day moving average (log scale)

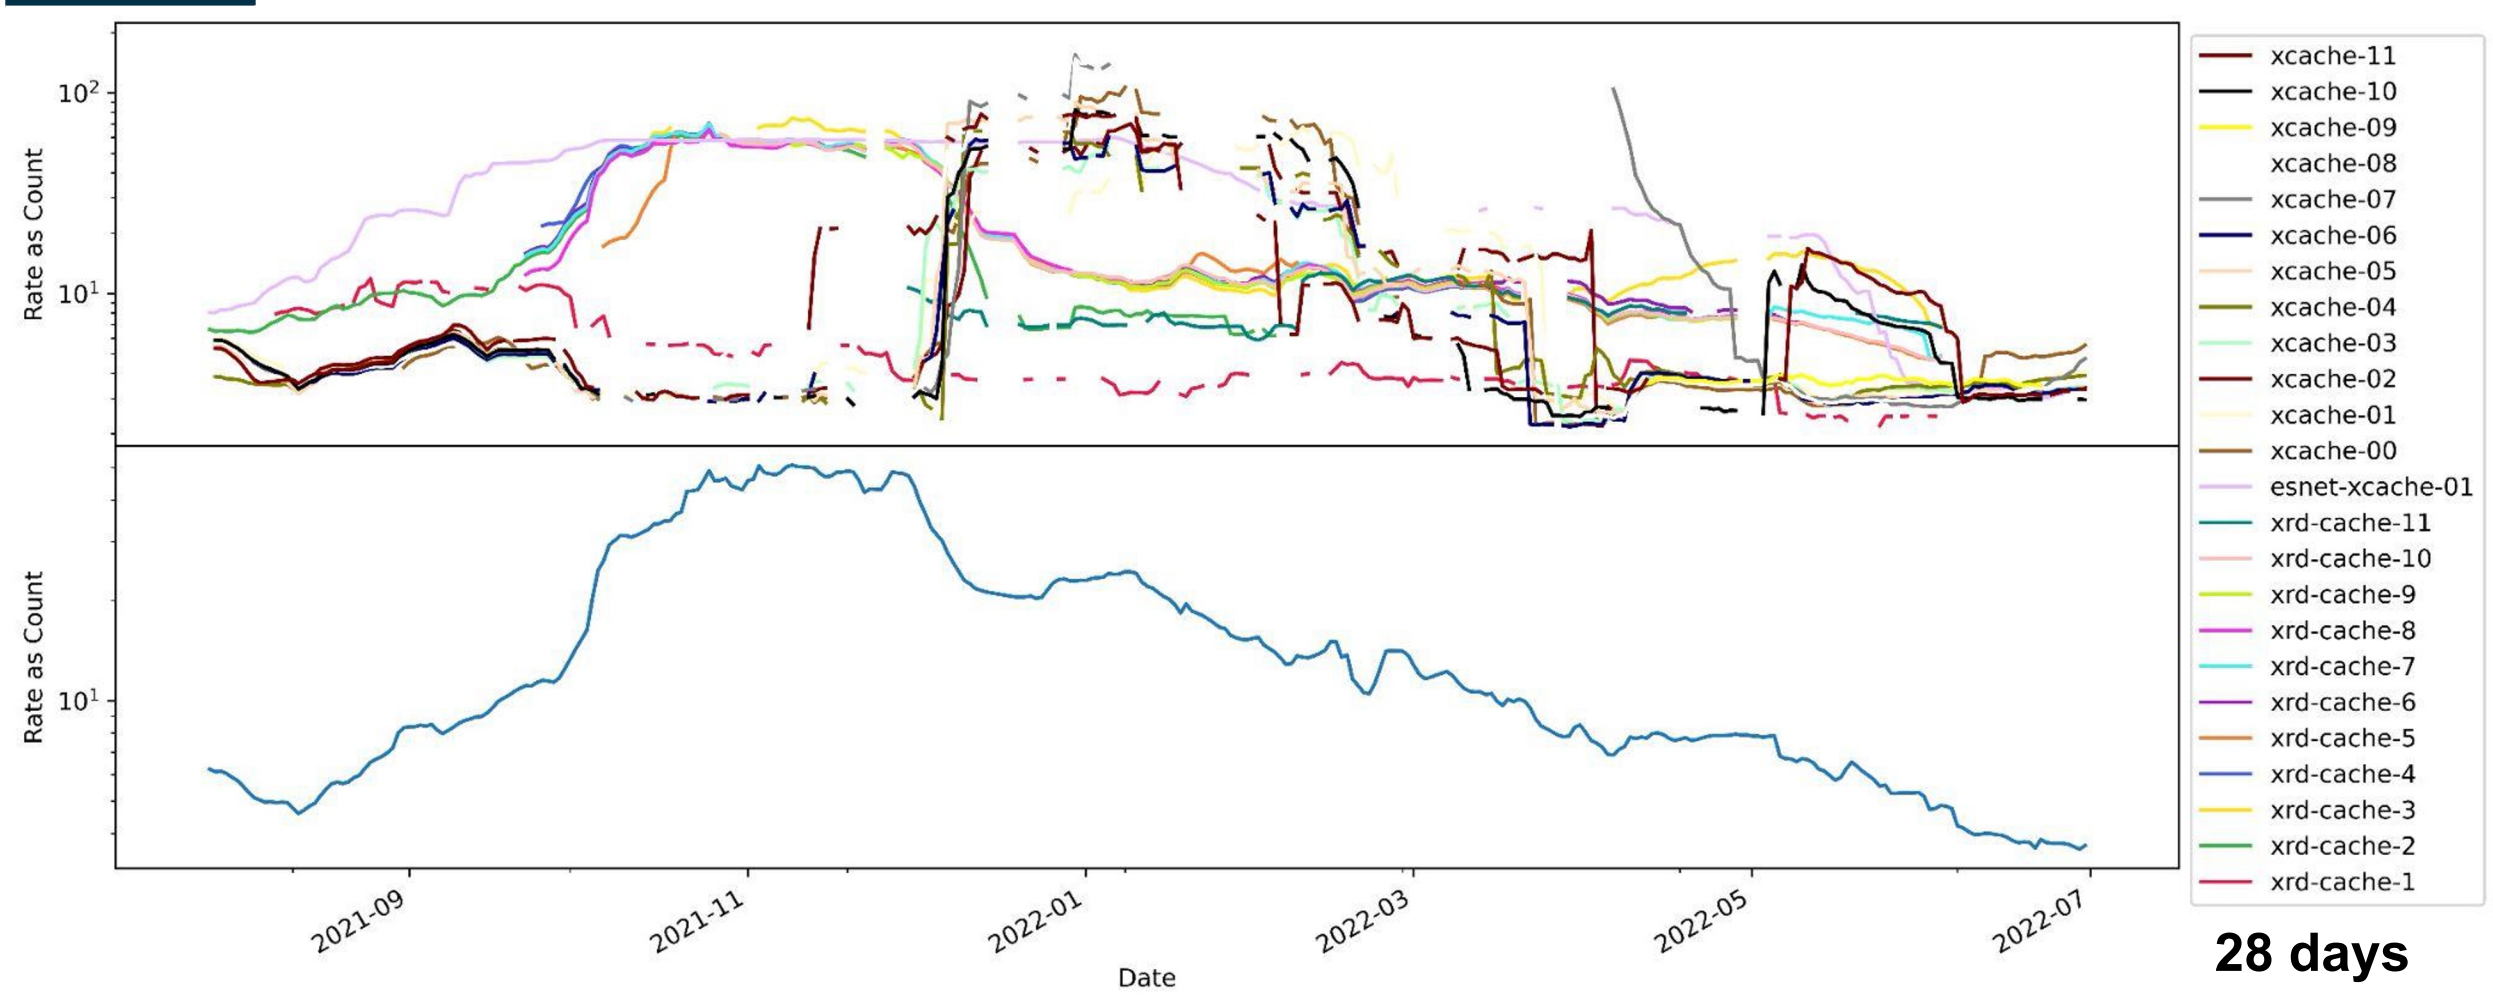

28 days

Daily data re-use volume

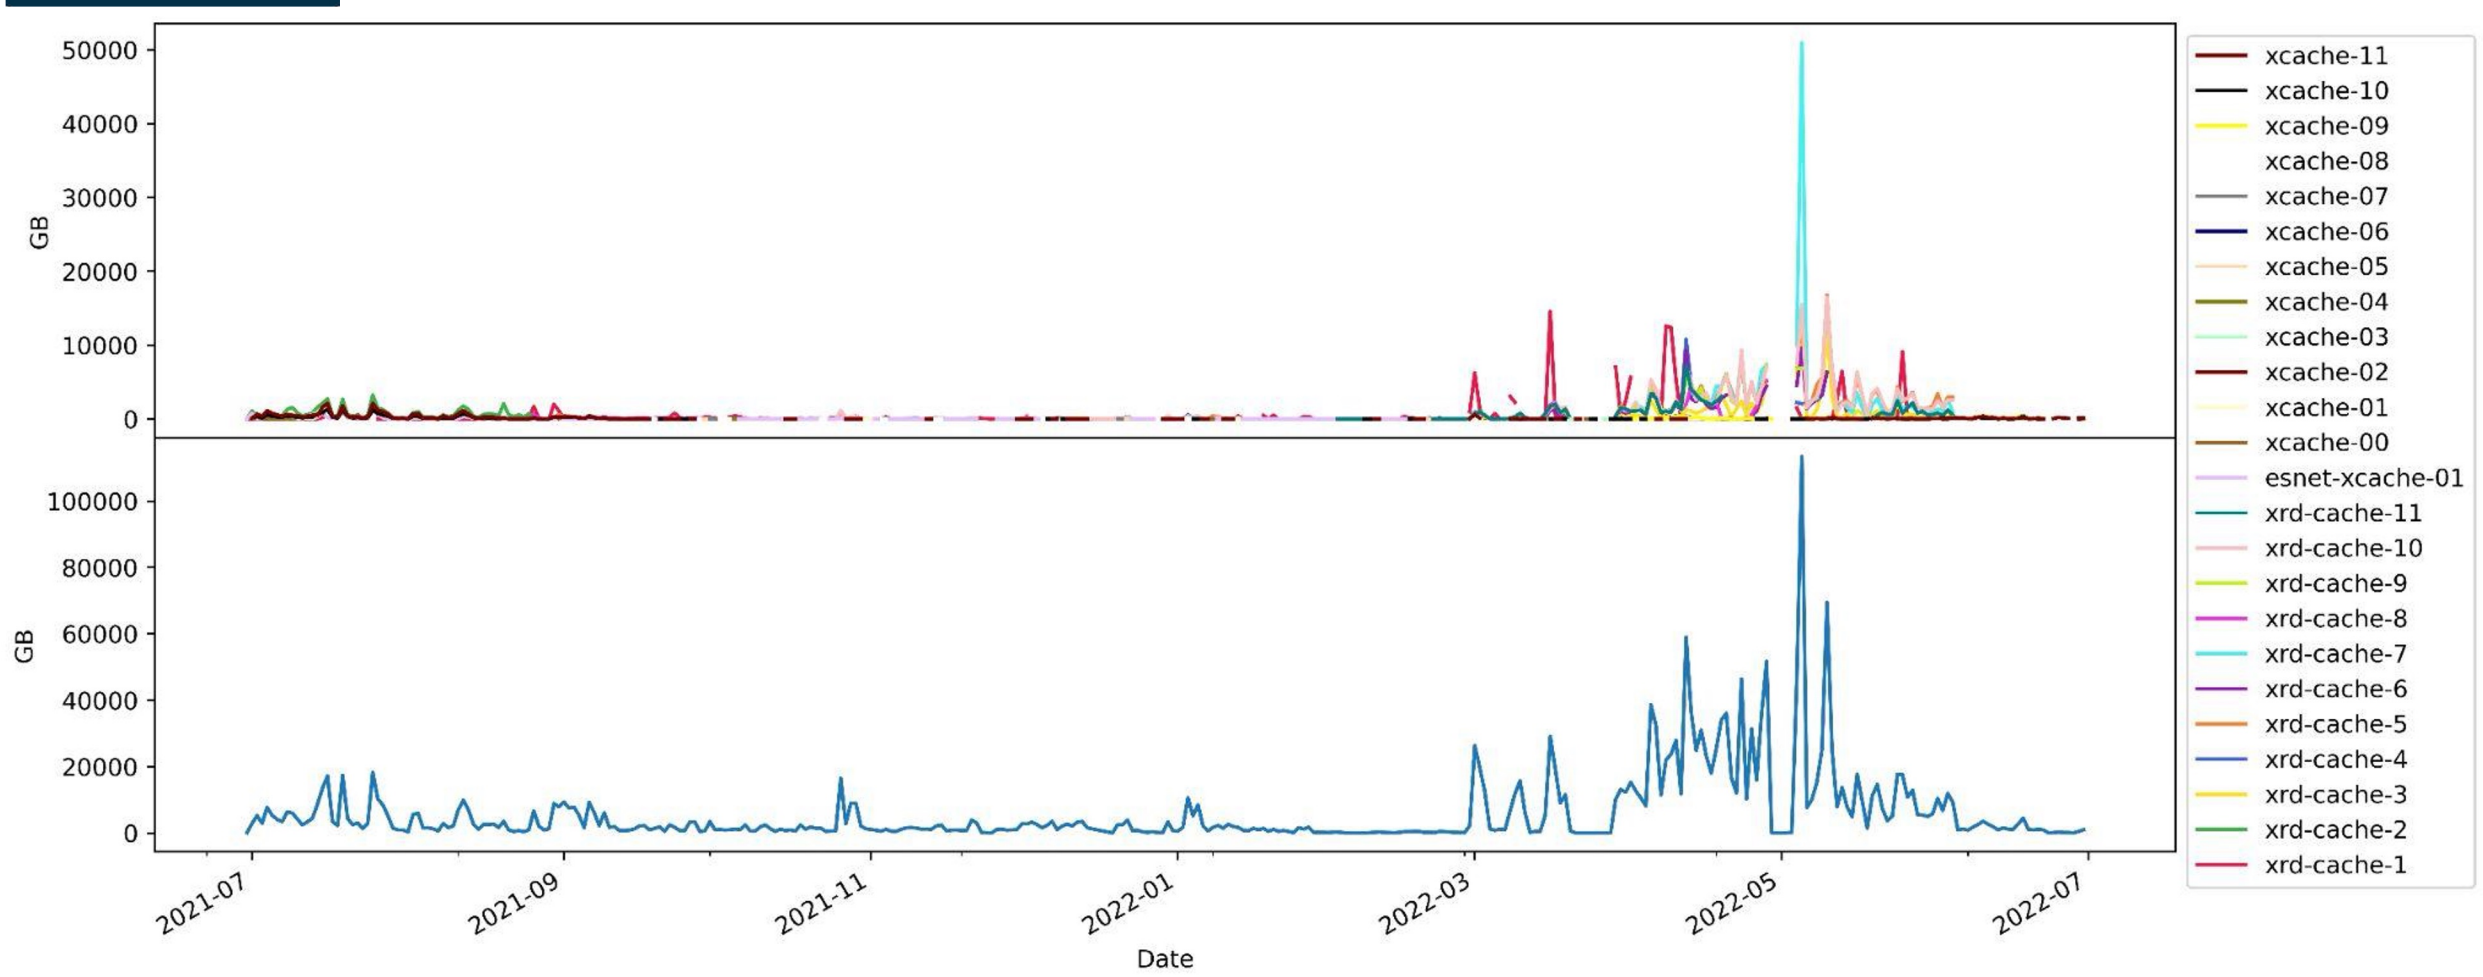

Daily data re-use volume with 2-day moving average

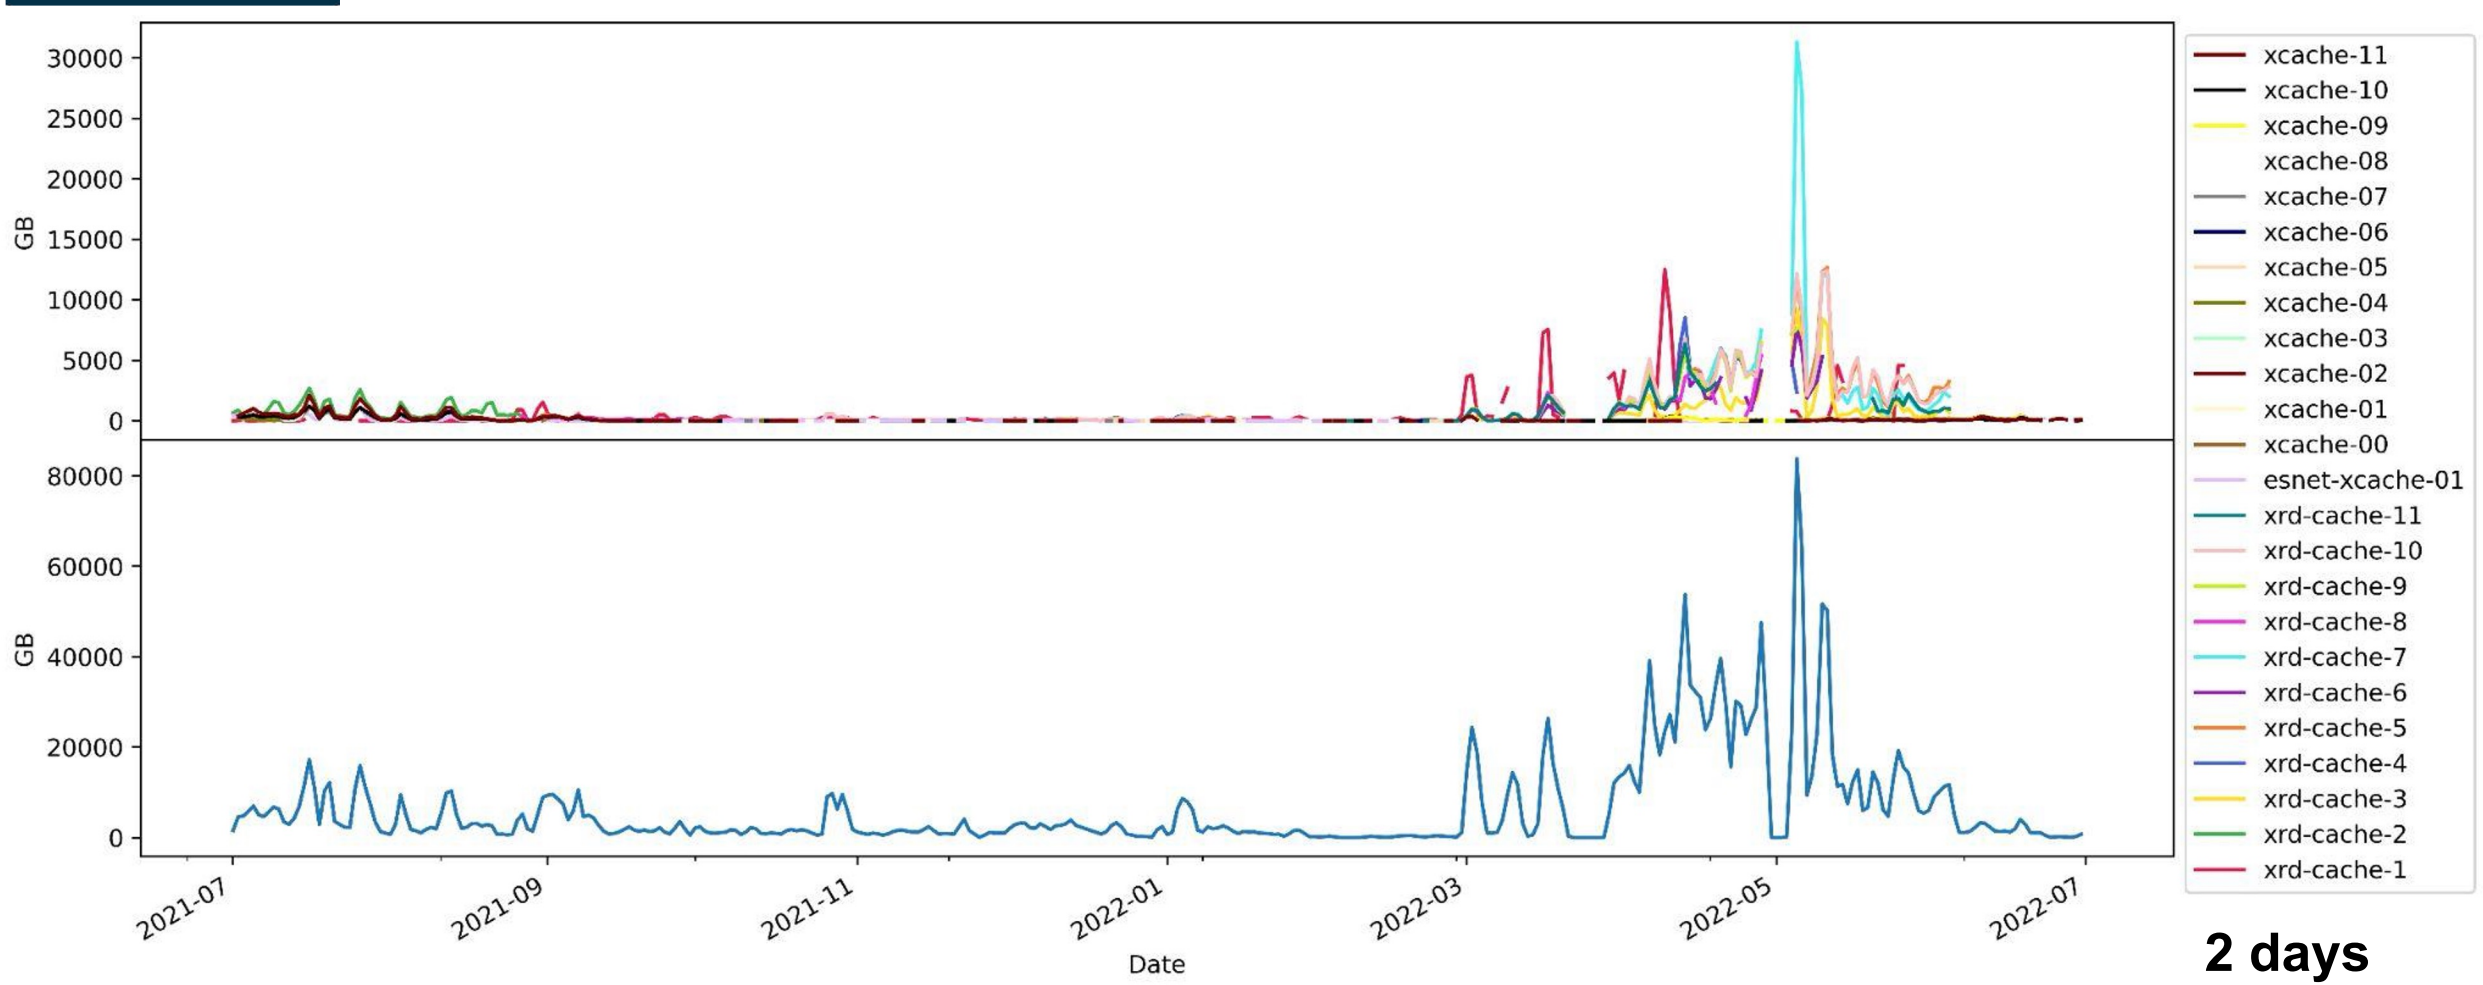

2 days

Daily data re-use volume with 3-day moving average

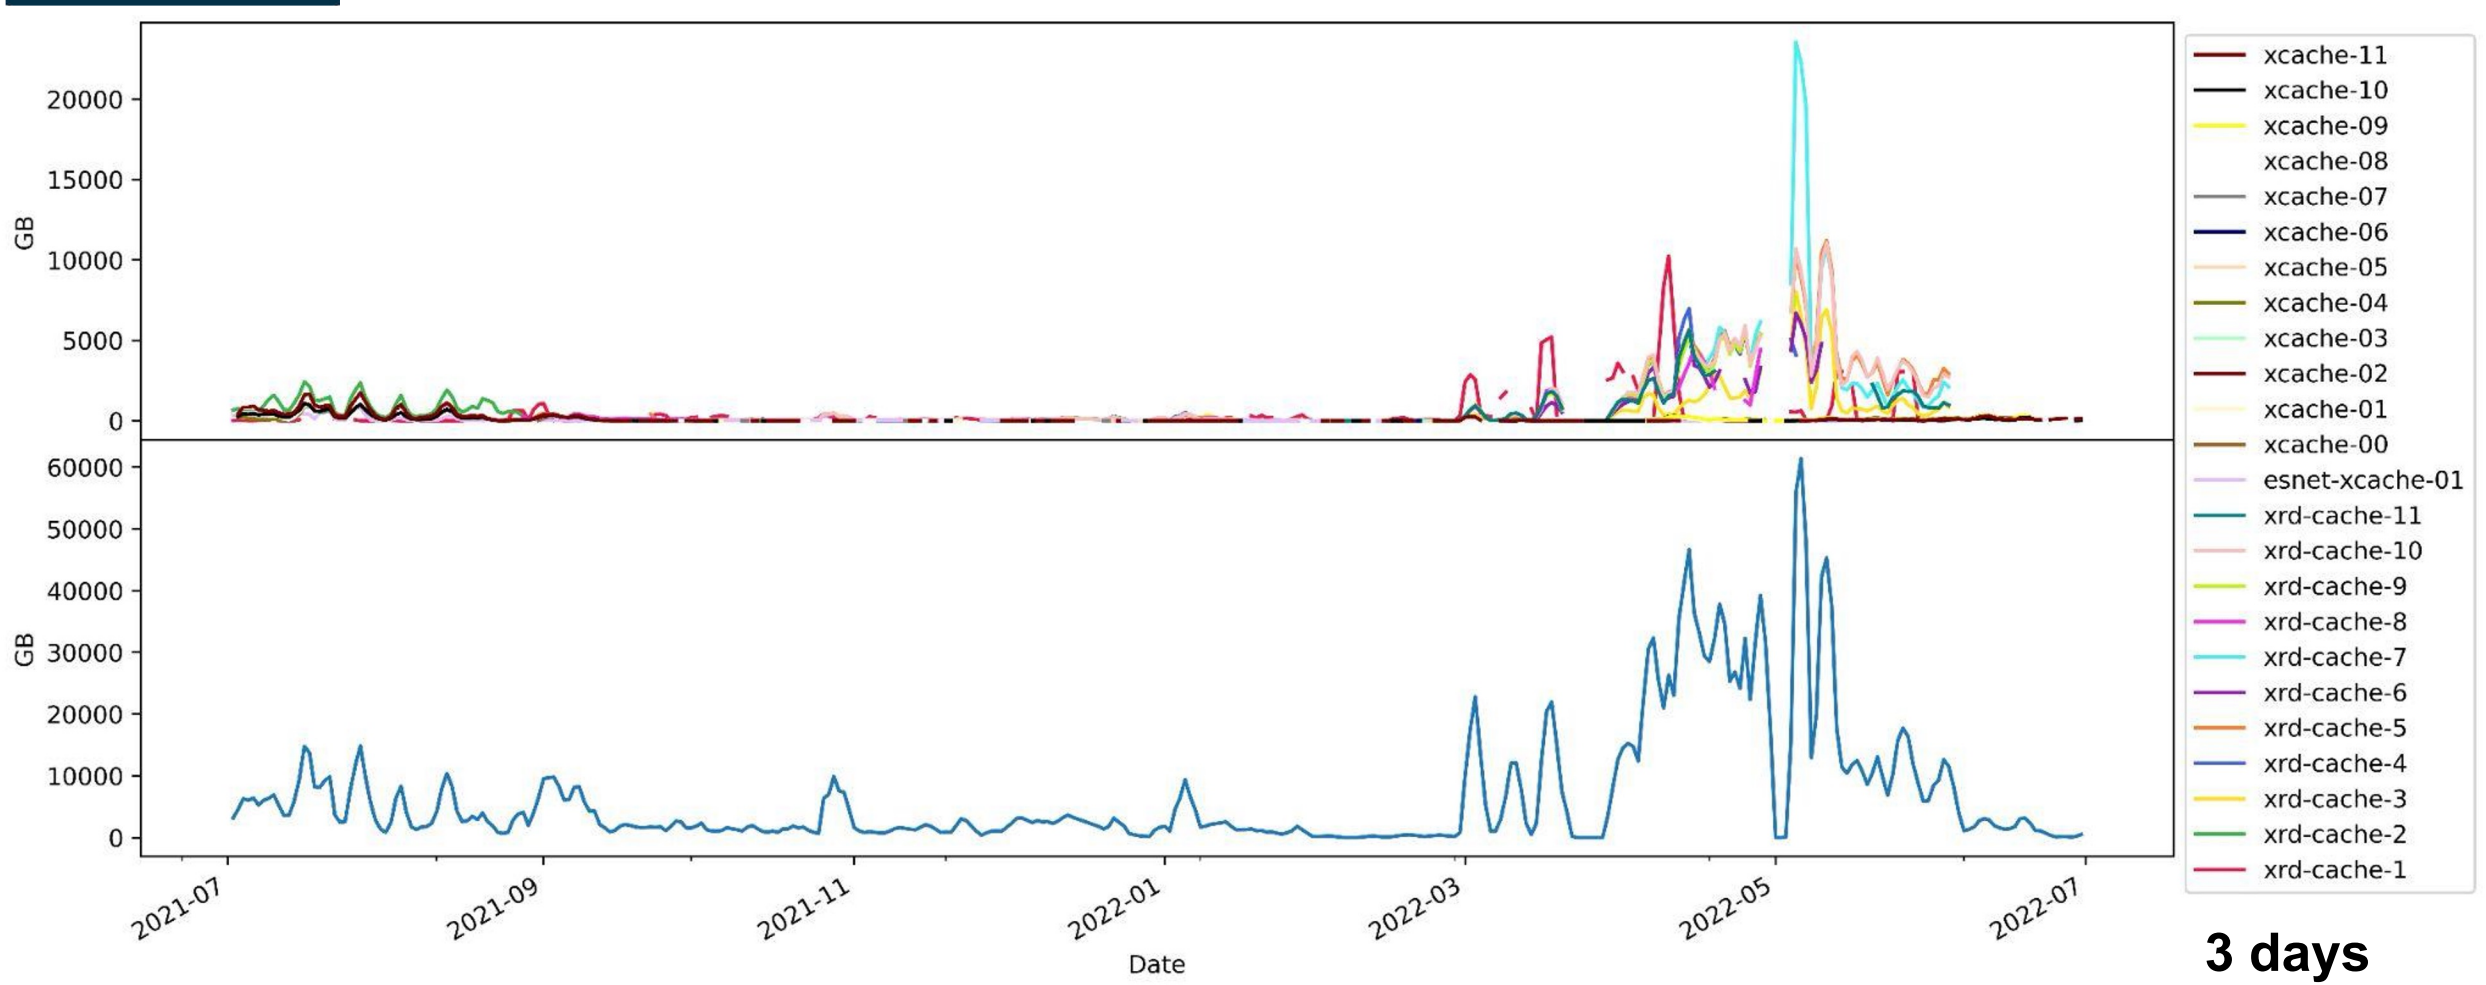

3 days

Daily data re-use volume with 5-day moving average

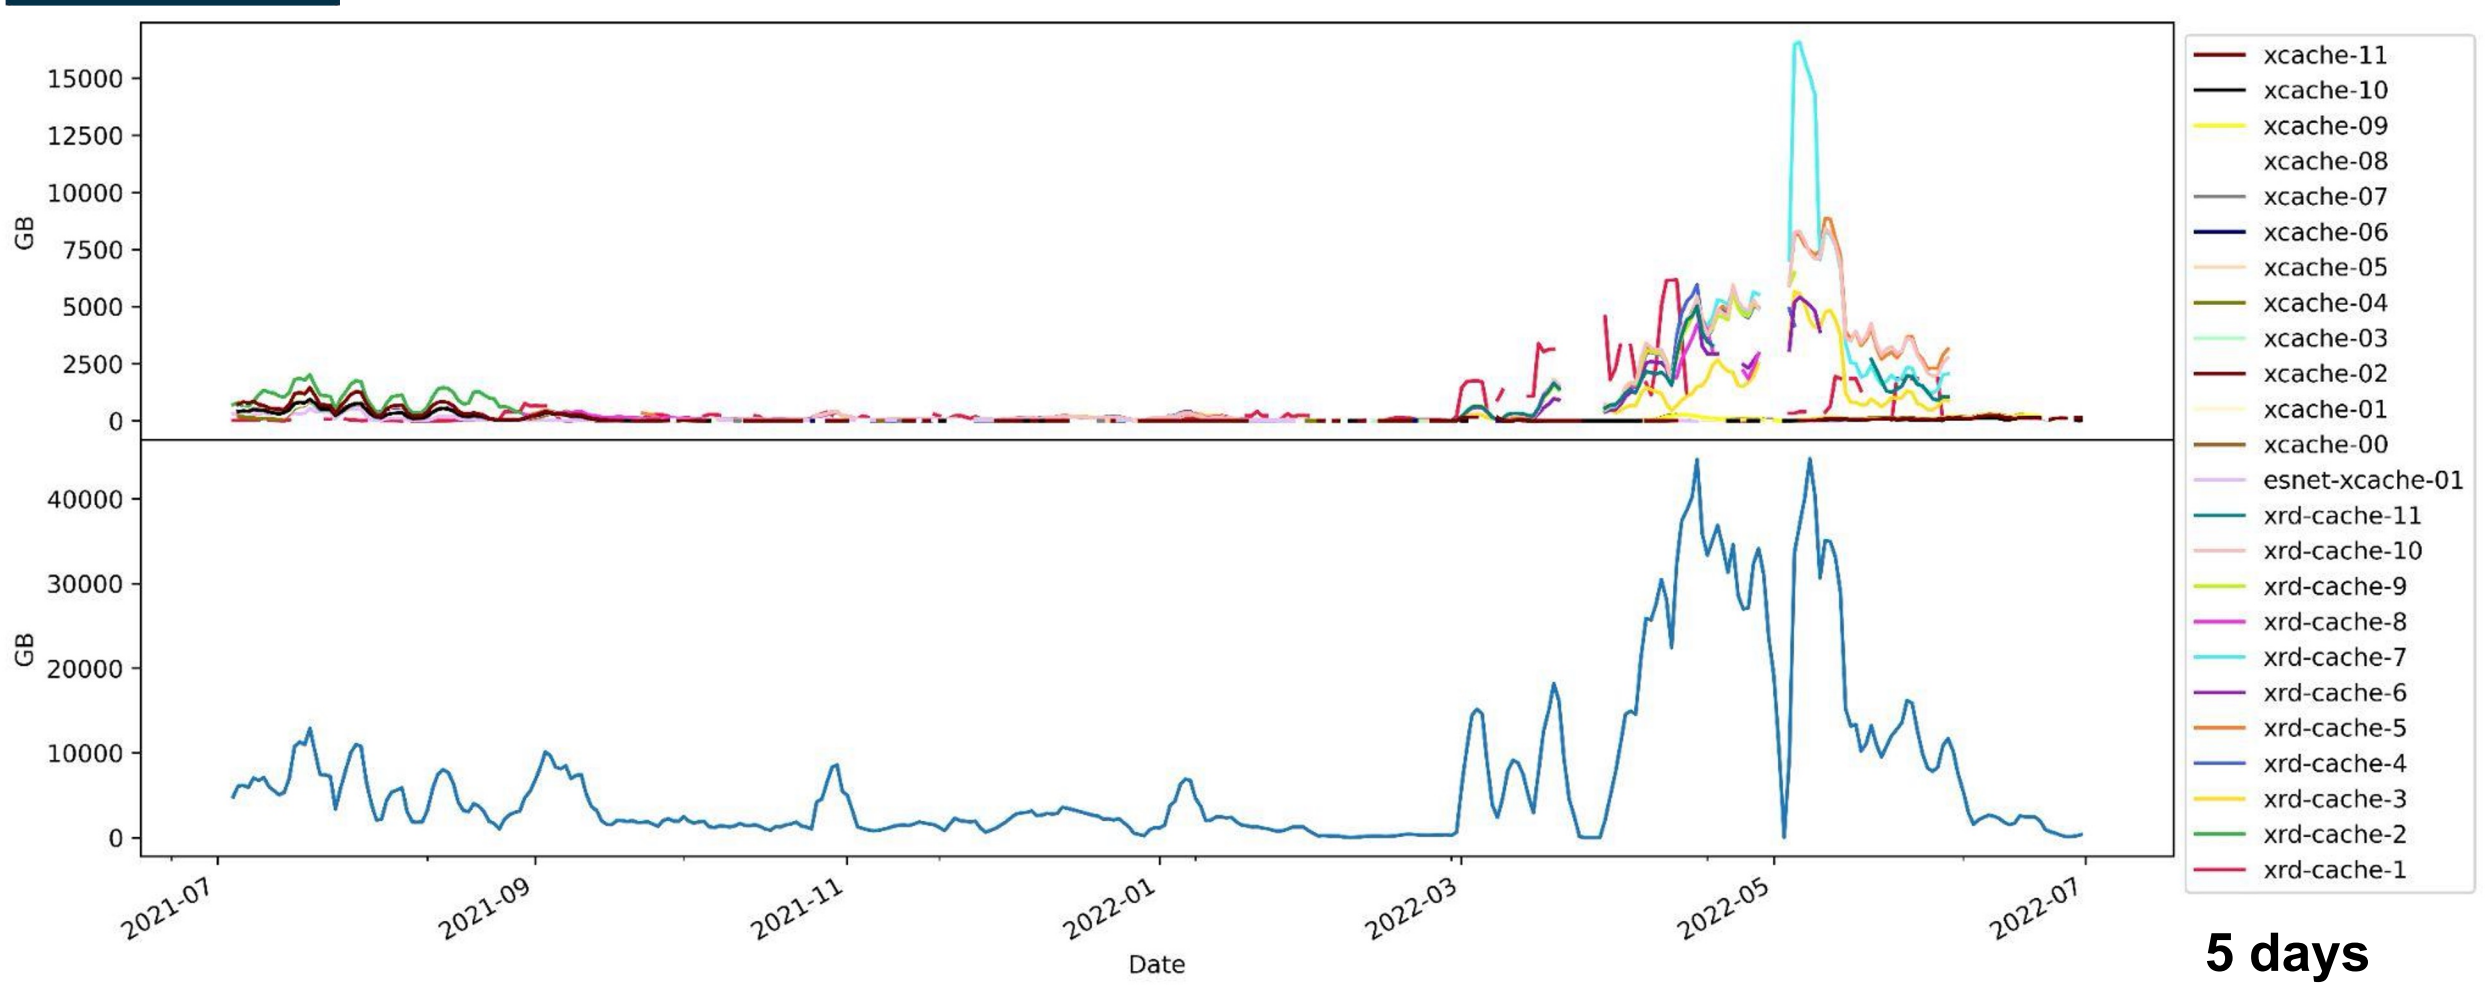

5 days

Daily data re-use volume with 7-day moving average

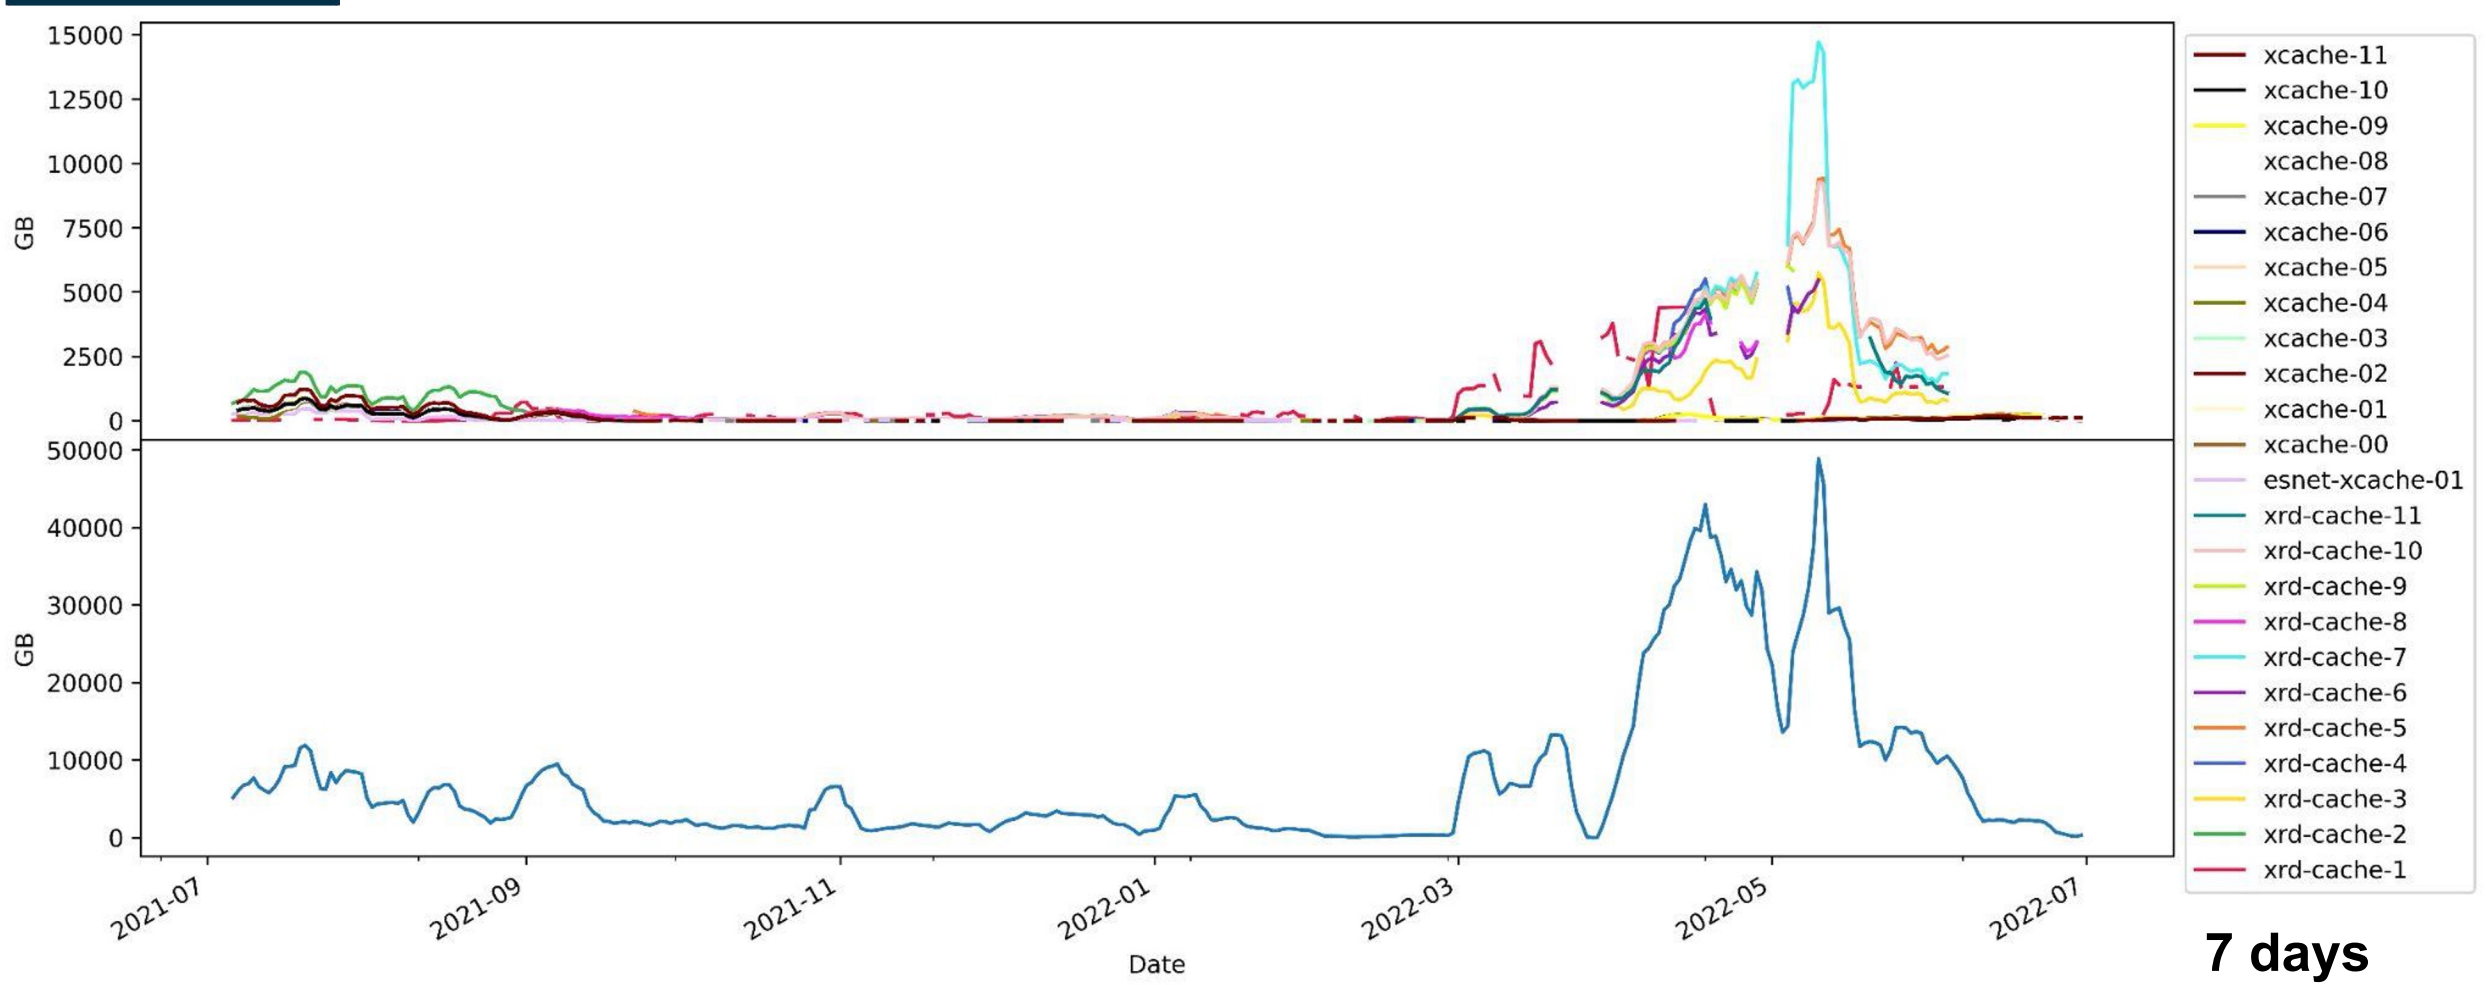

7 days

Daily data re-use volume with 10-day moving average

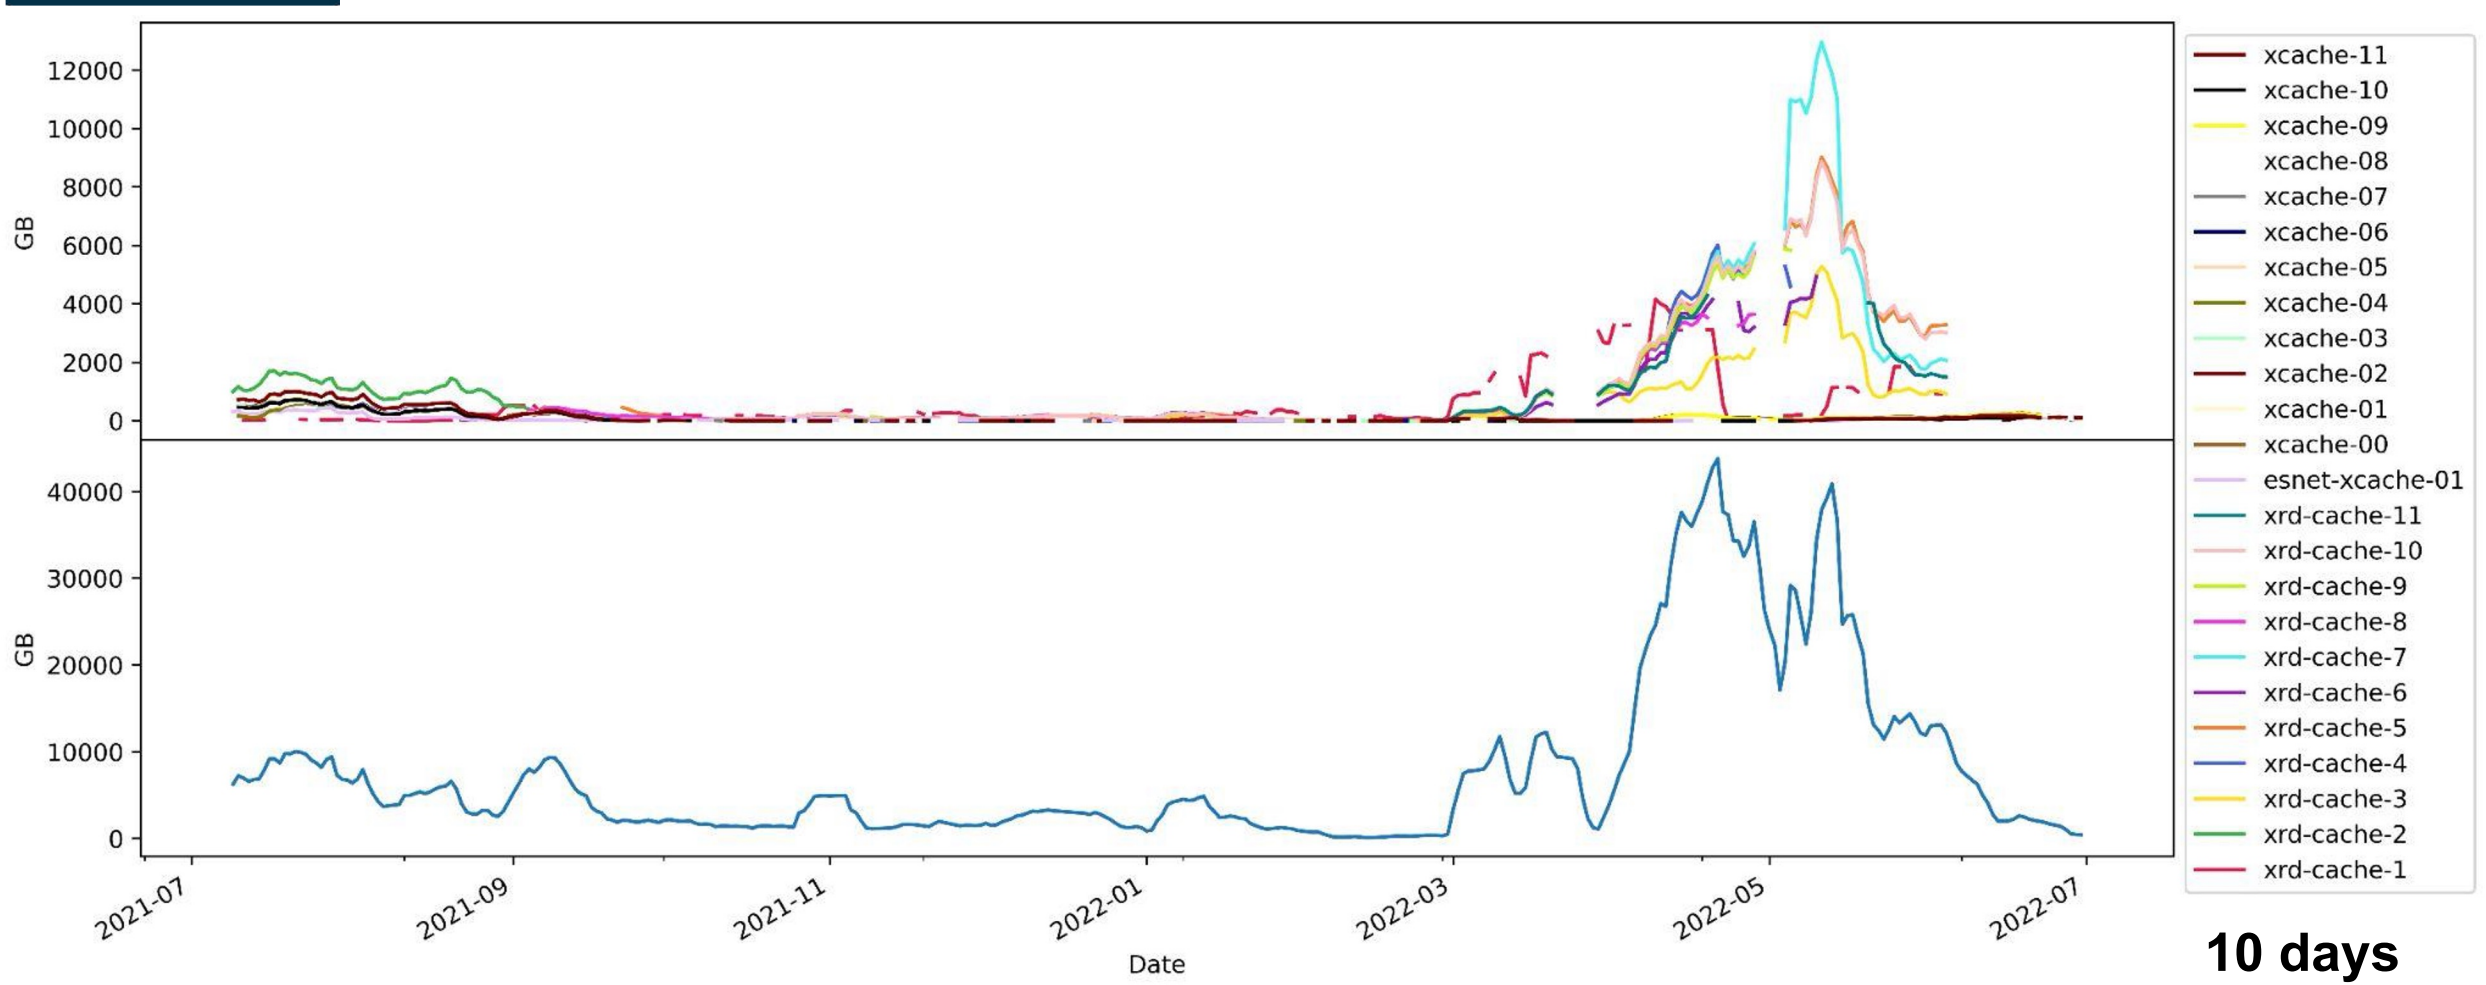

10 days

Daily data re-use volume with 14-day moving average

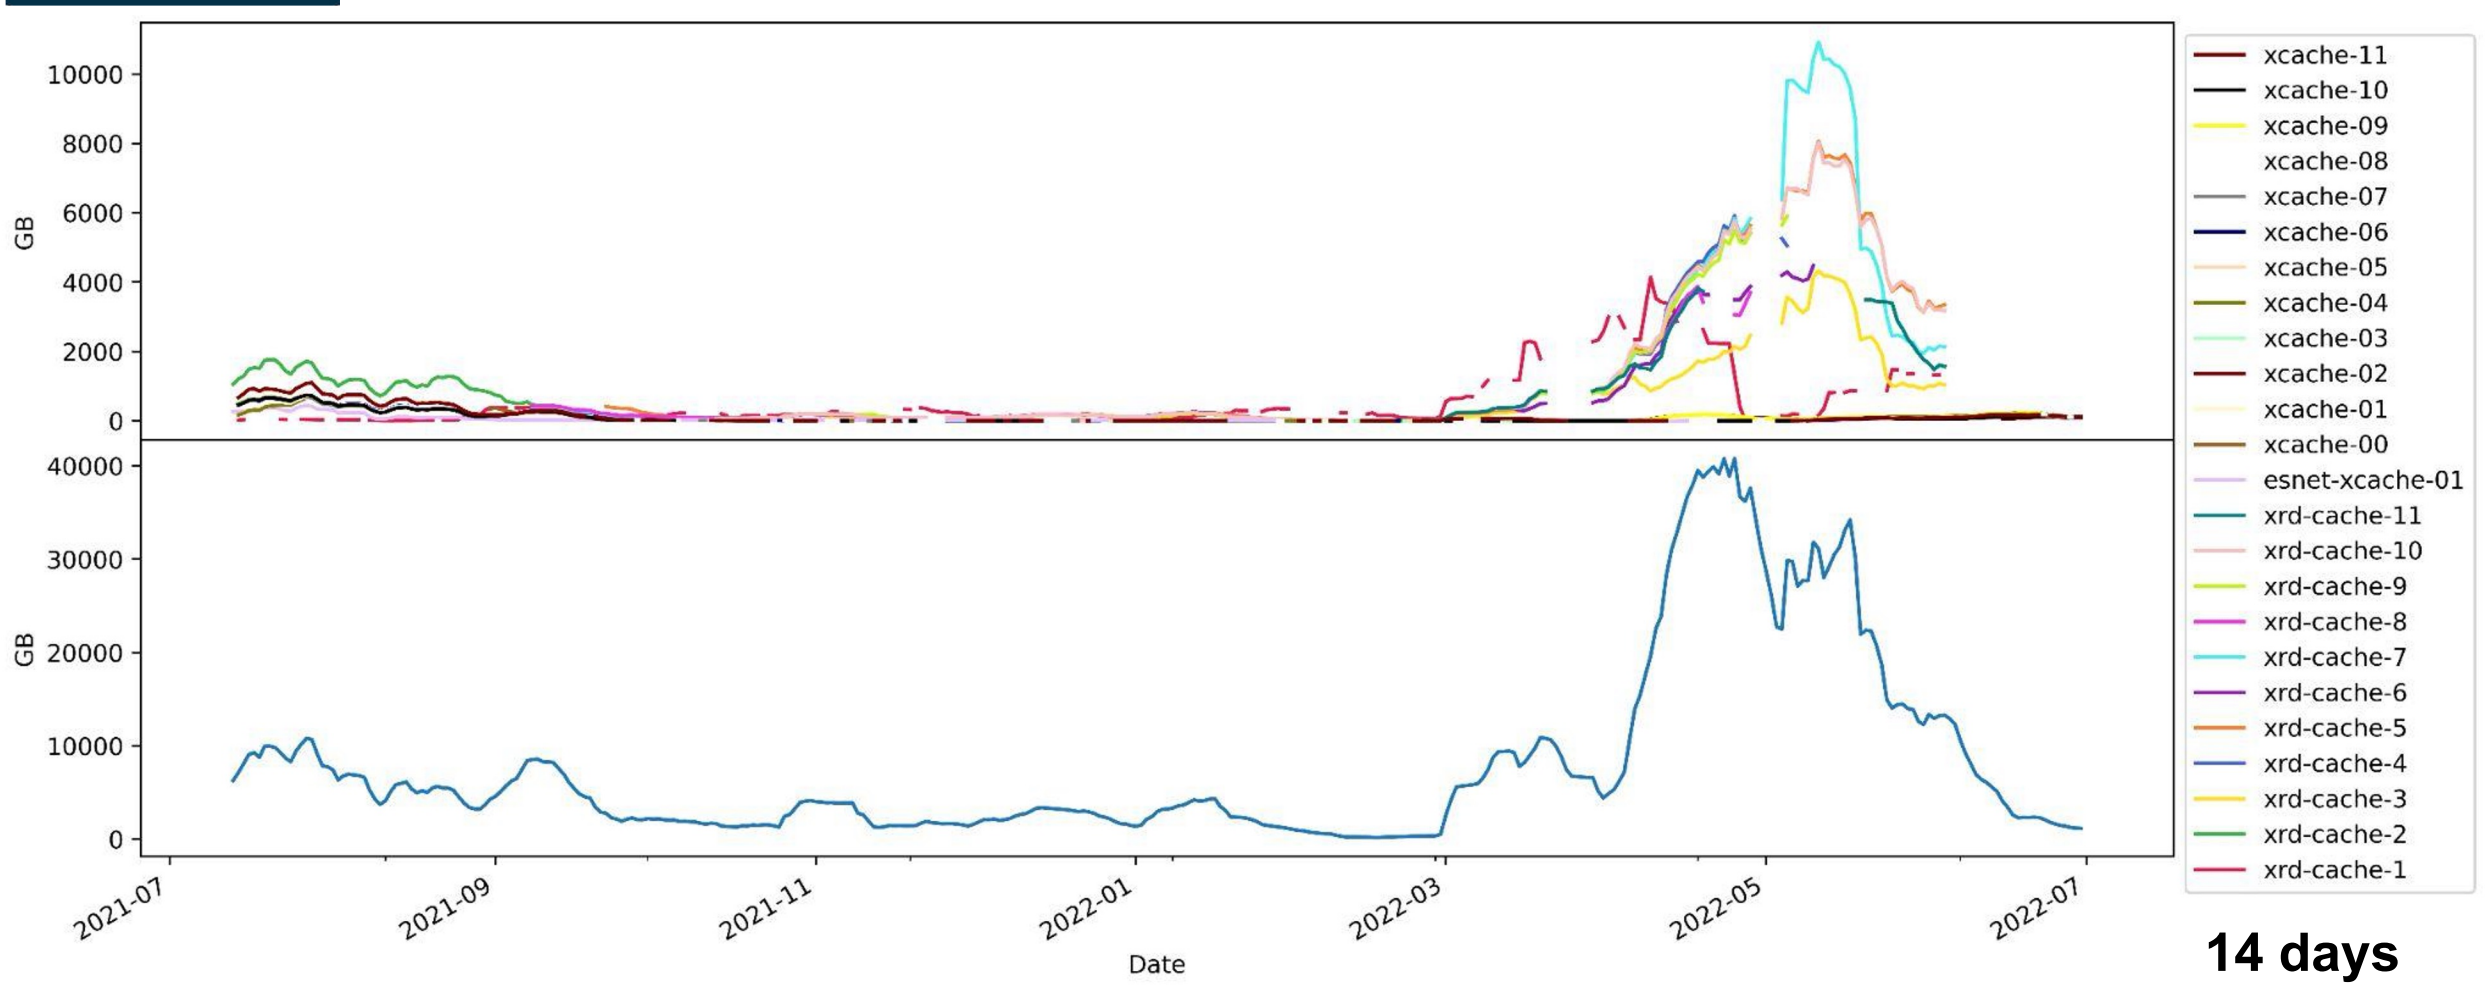

14 days

Daily data re-use volume with 21-day moving average

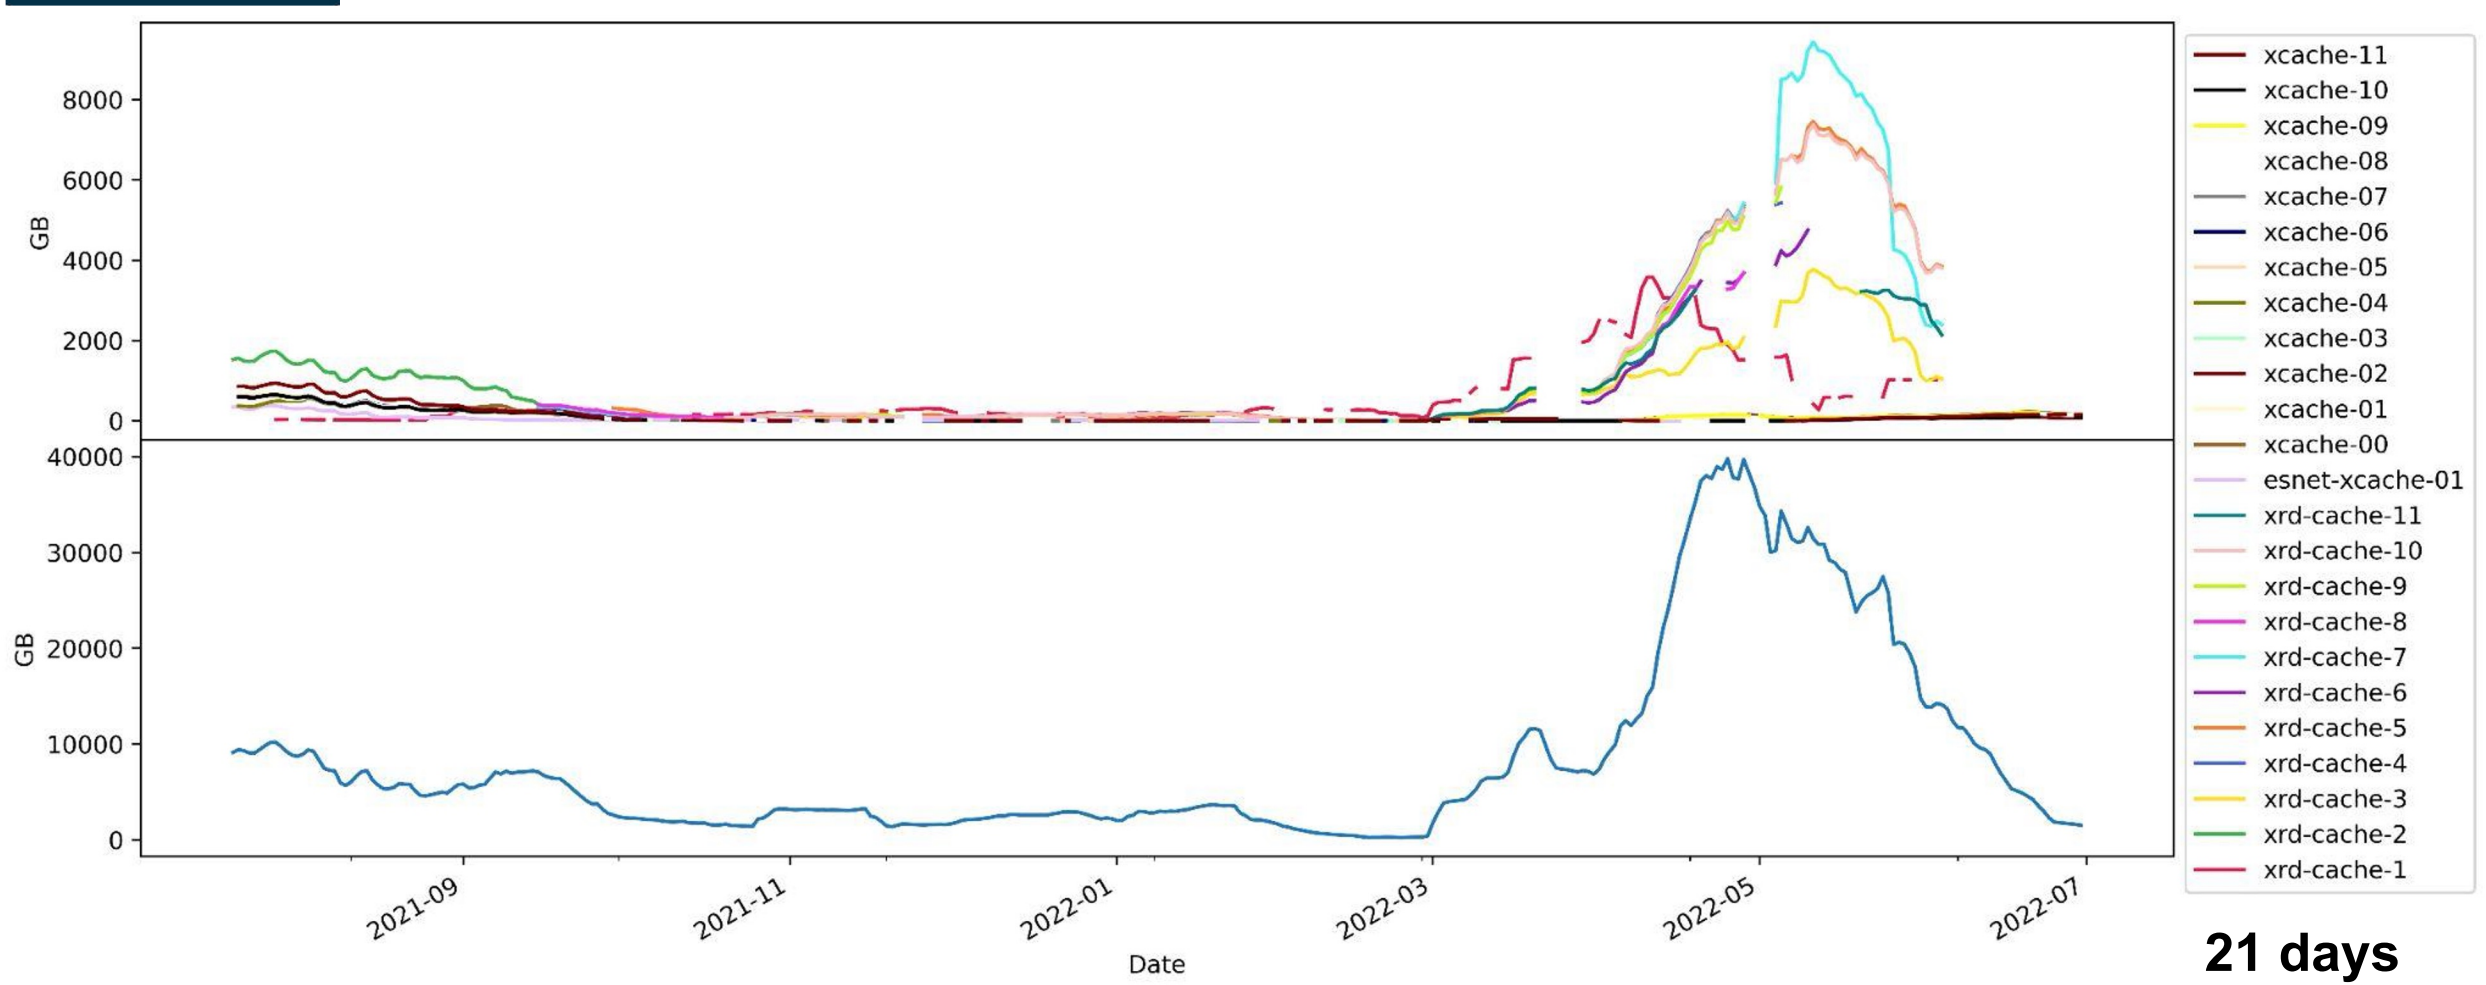

21 days

Daily data re-use volume with 28-day moving average

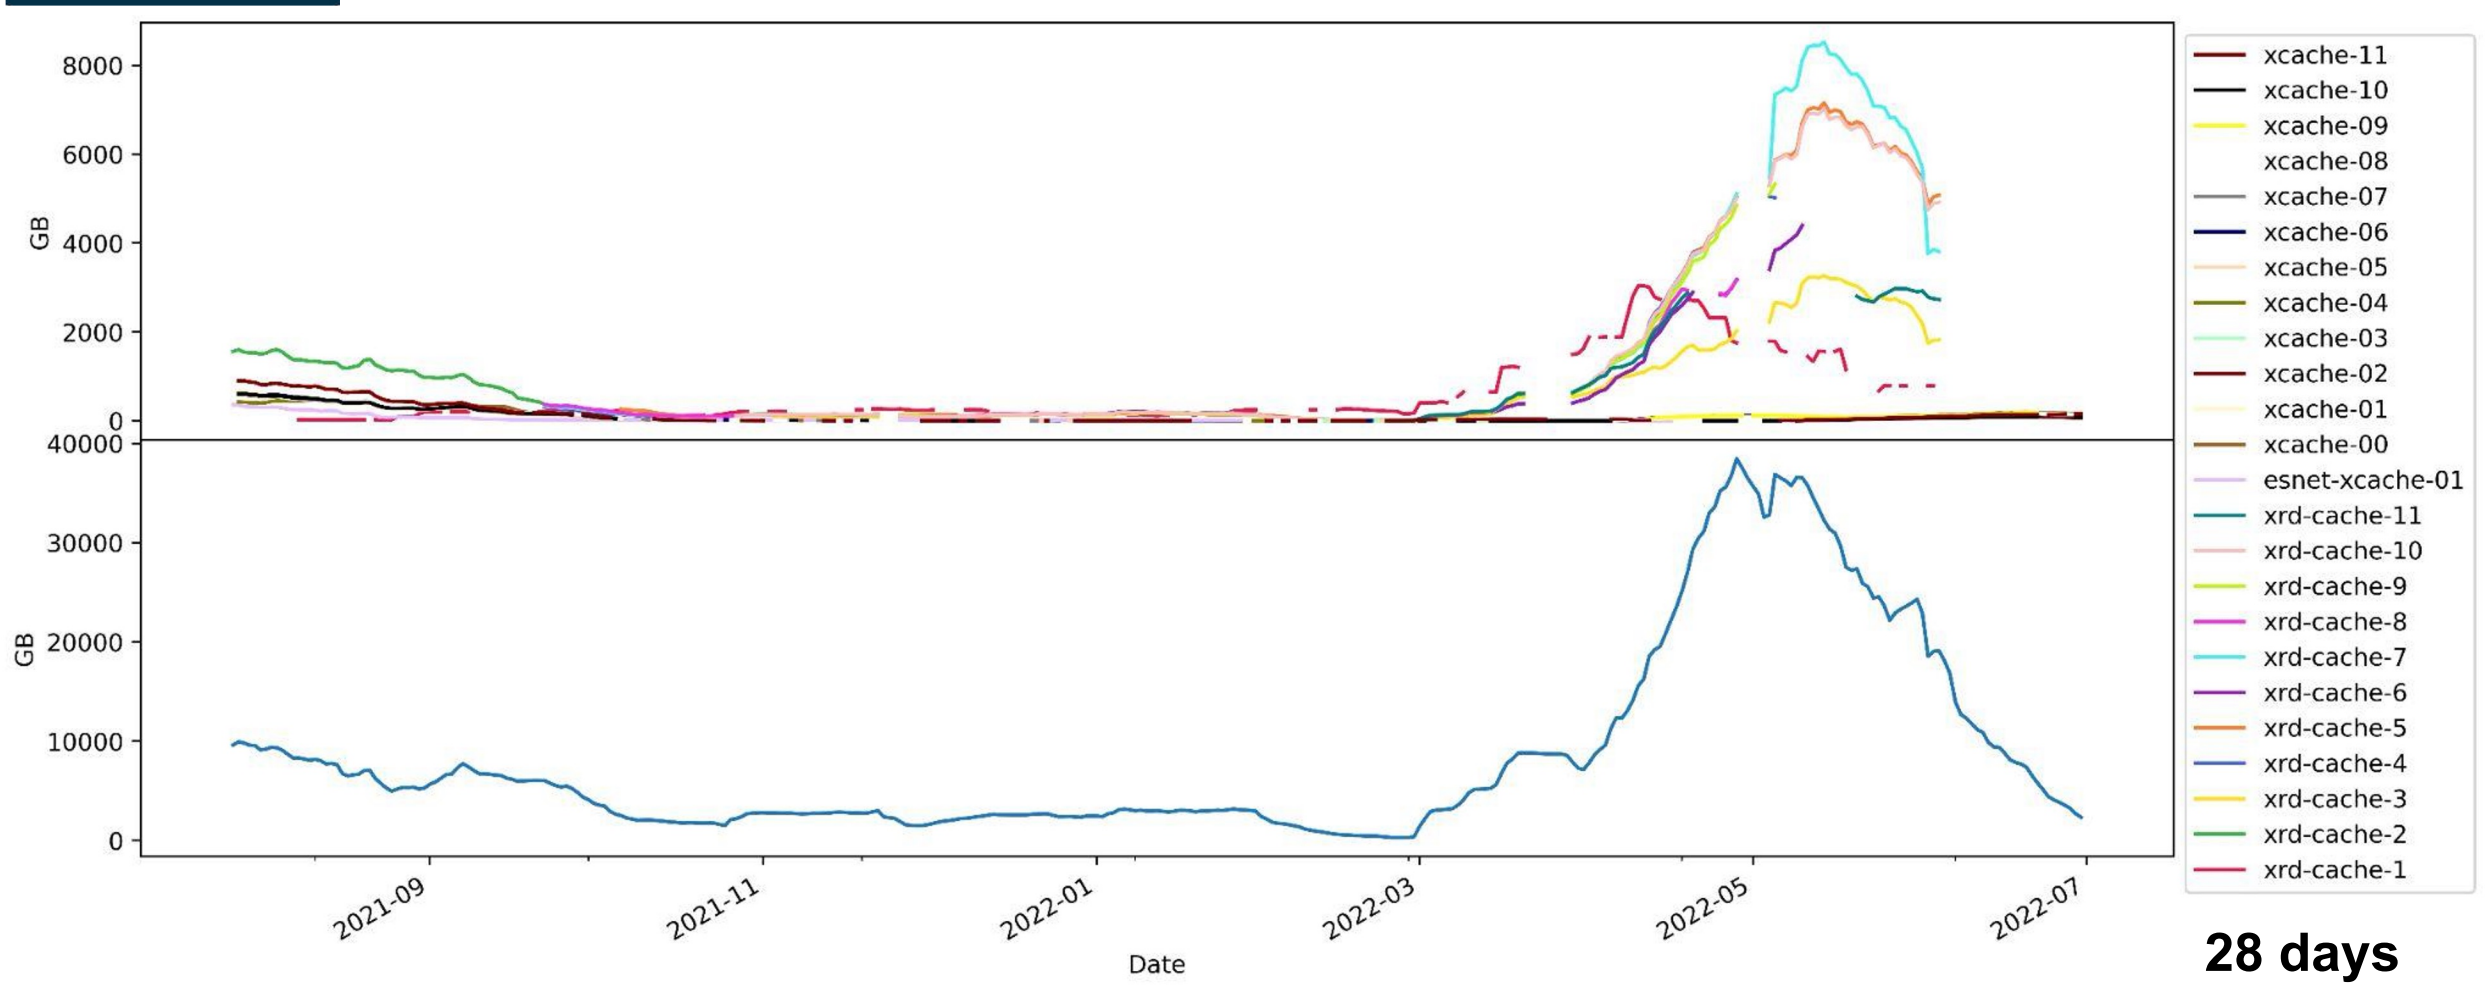

28 days

Daily data re-access rates (log scale)

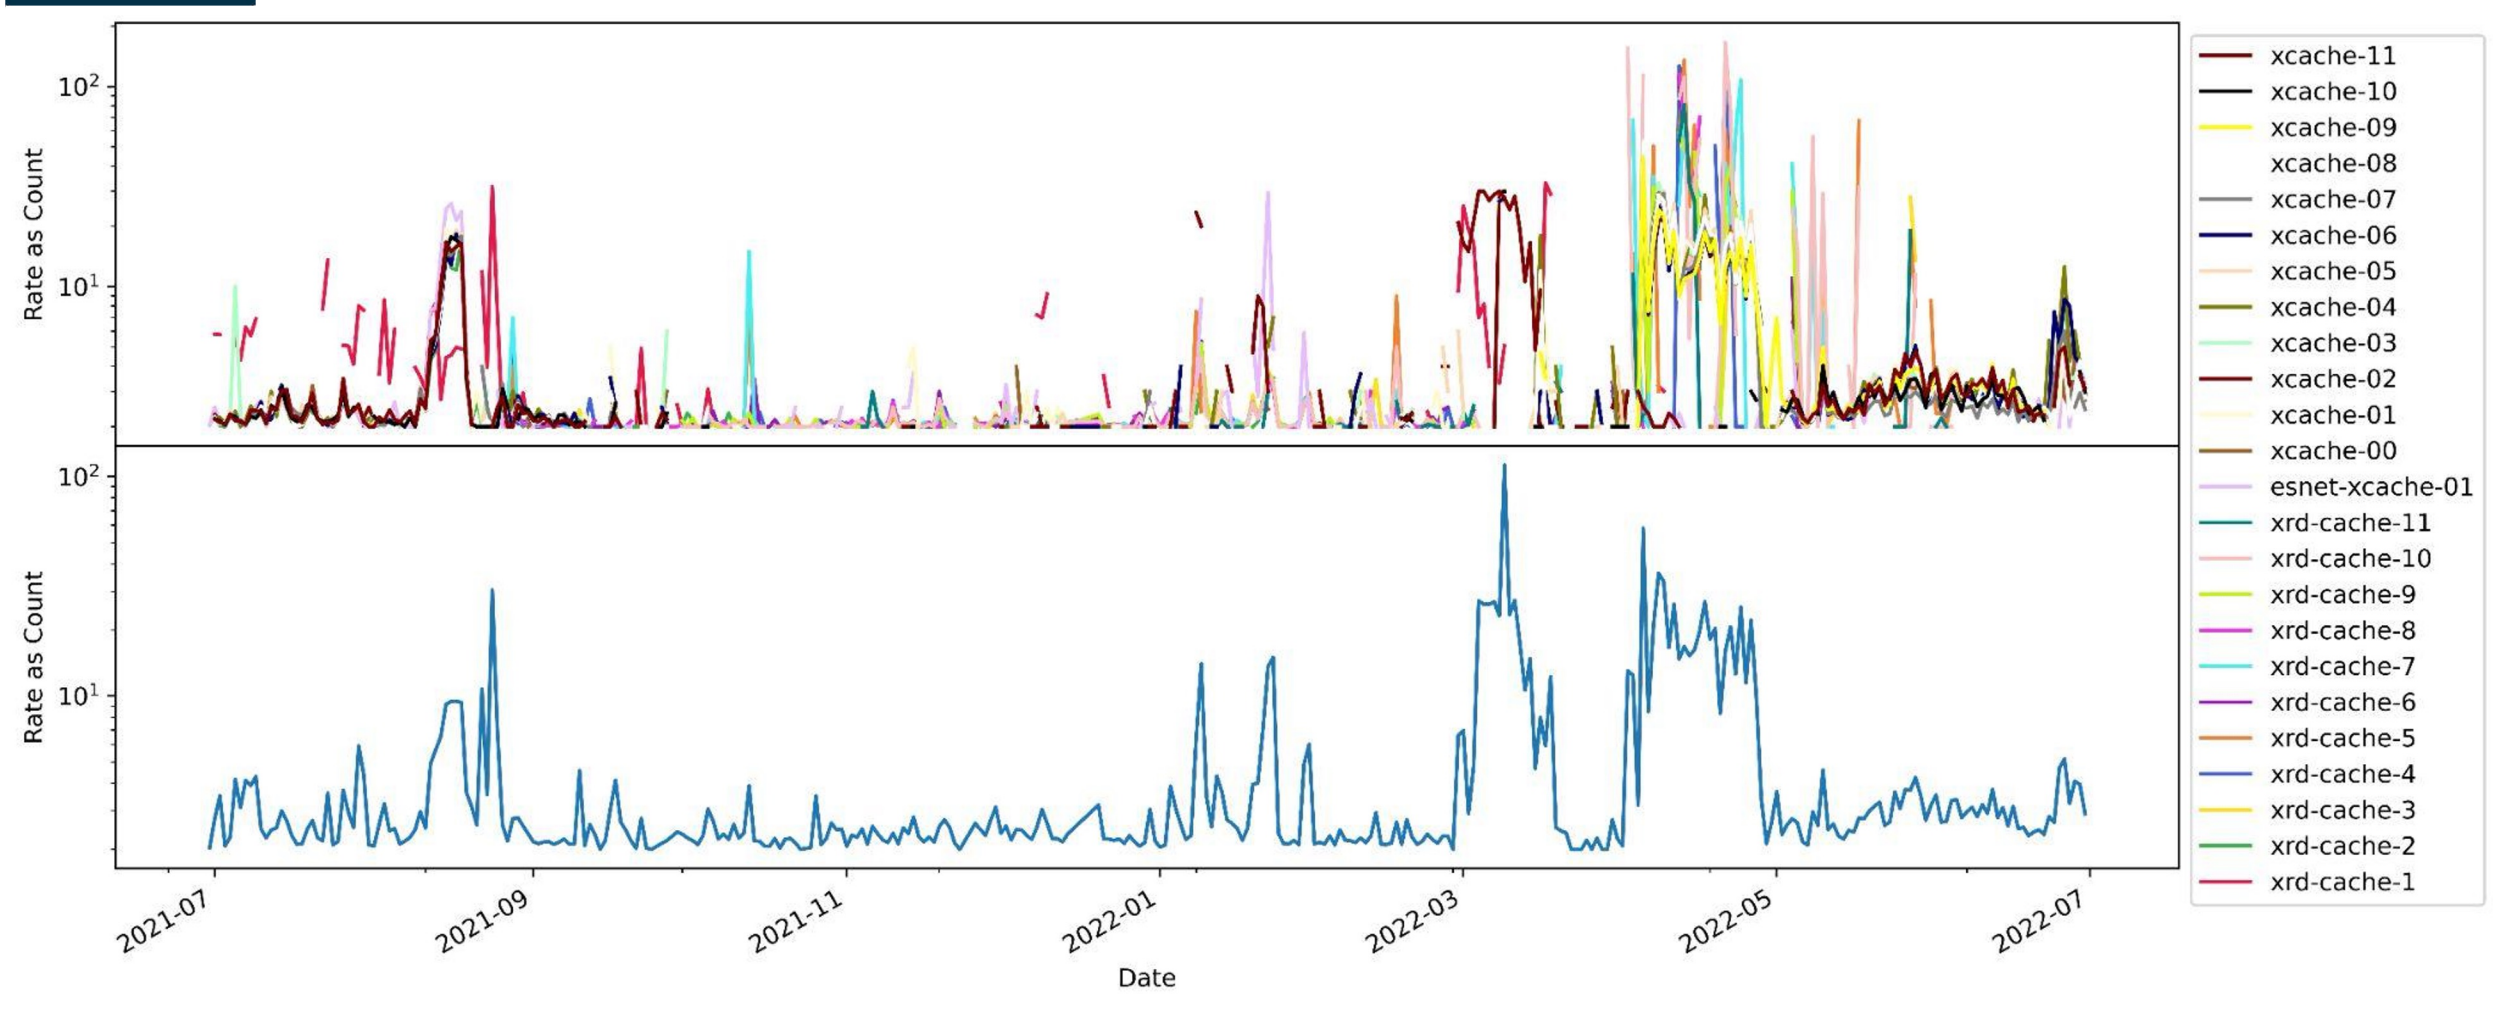

Daily data re-access rates with moving average 2 days (log scale)

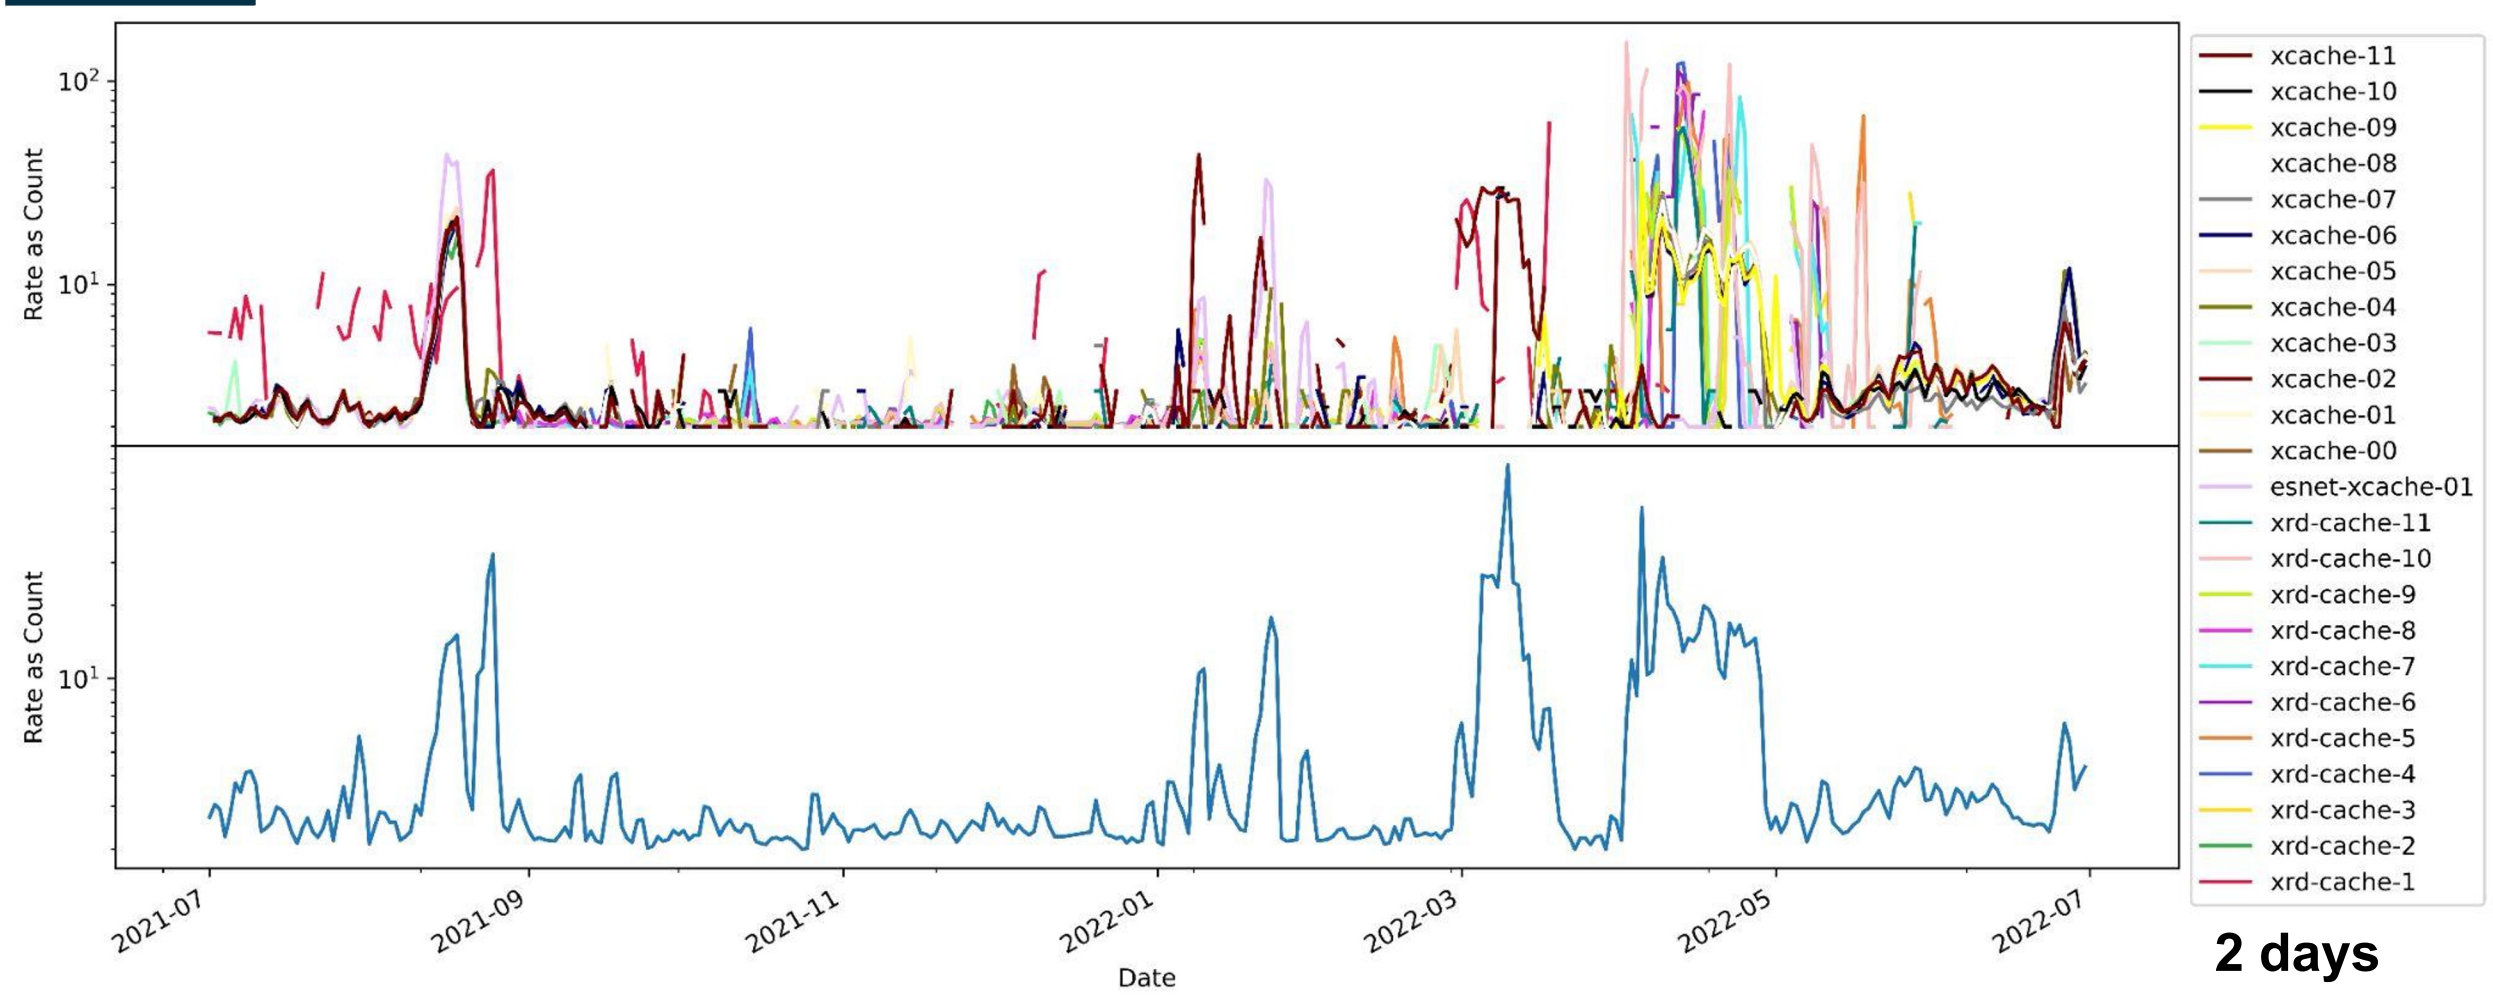

2 days

Daily data re-access rates with moving average 3 days (log scale)

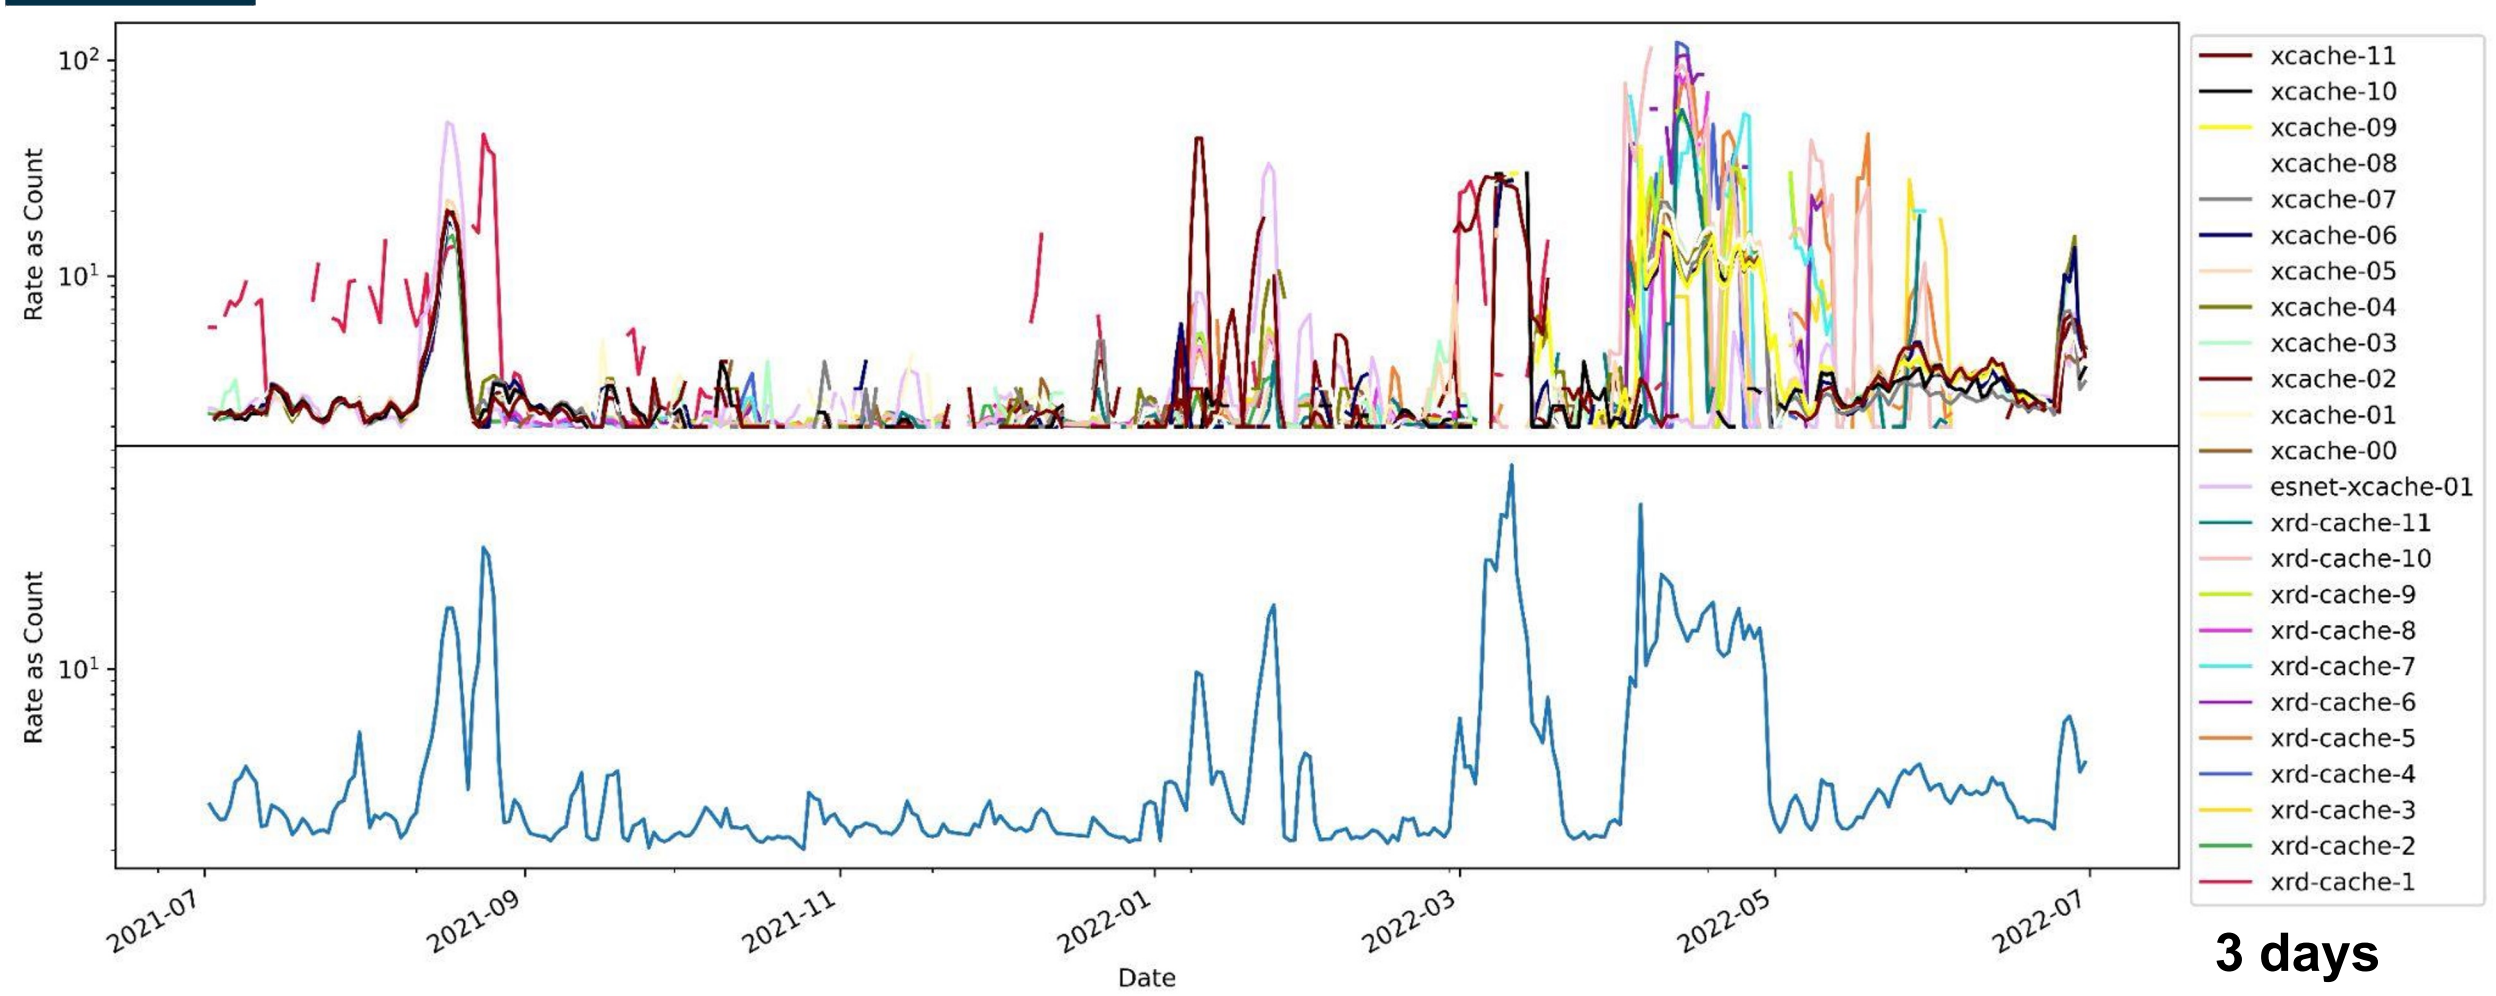

3 days

Daily data re-access rates with moving average 5 days (log scale)

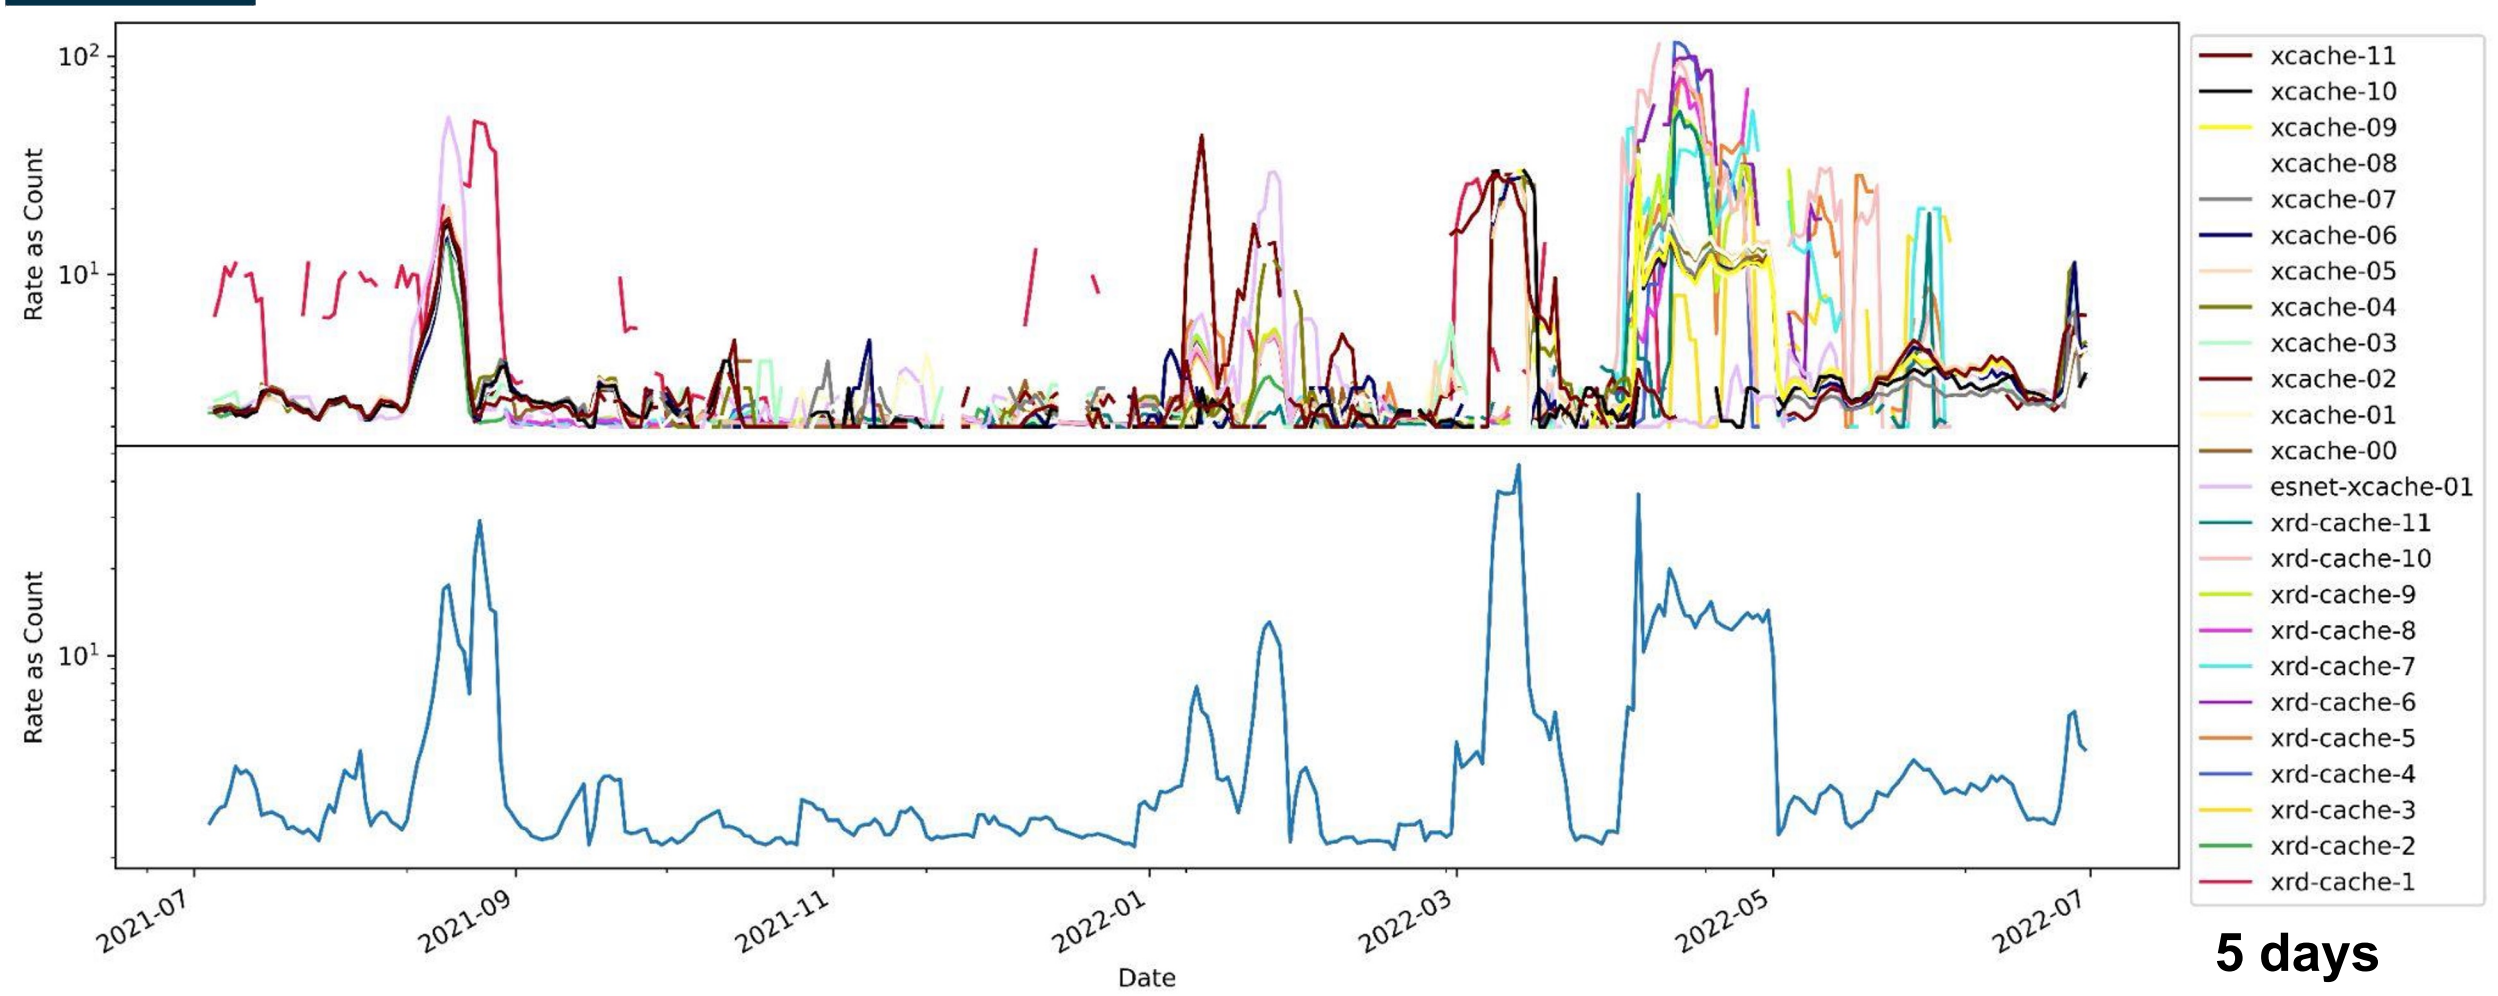

5 days

Daily data re-access rates with moving average 7 days (log scale)

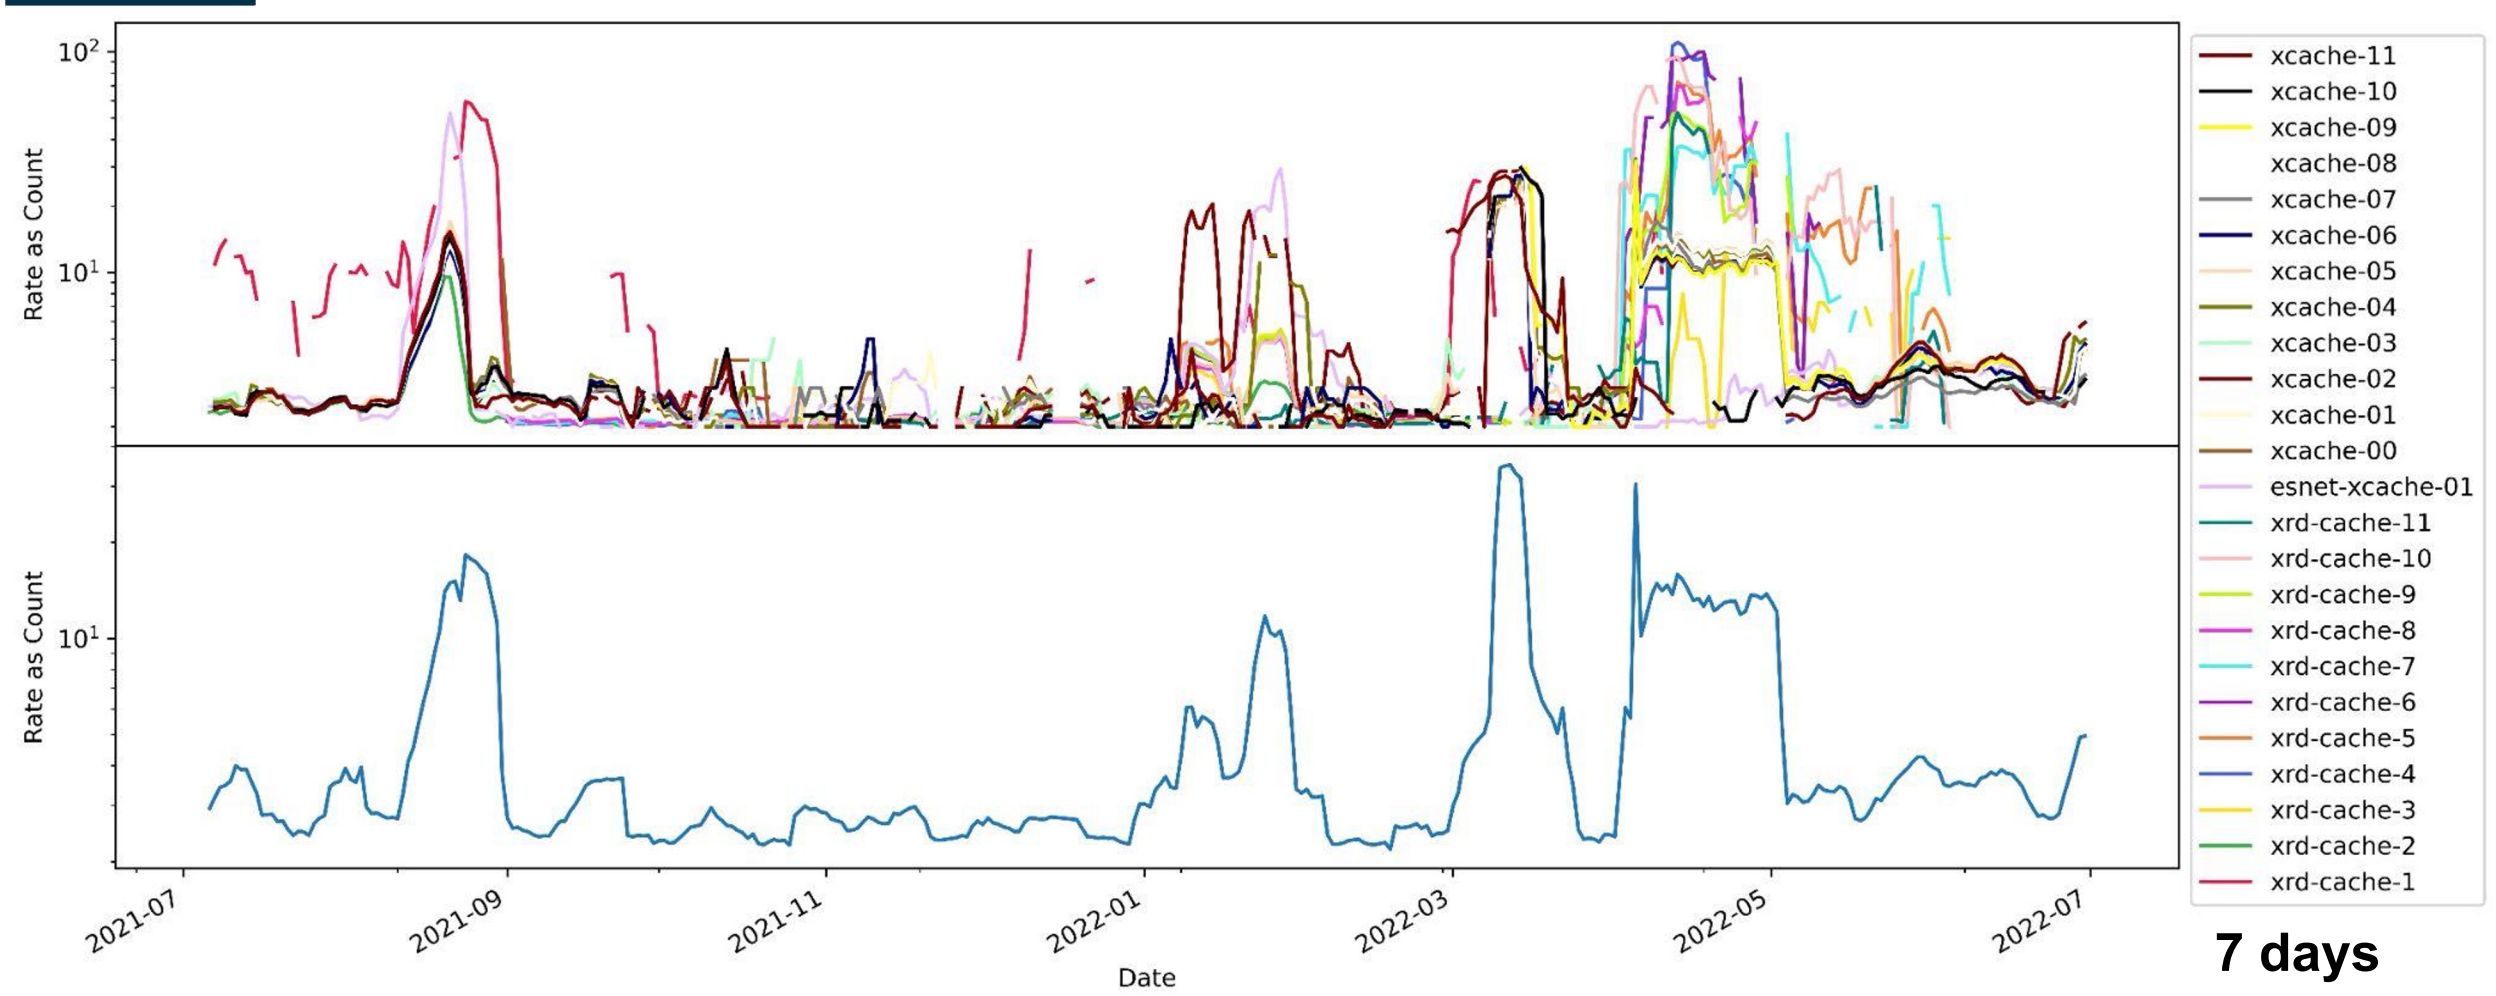

7 days

Daily data re-access rates with moving average 10 days (log scale)

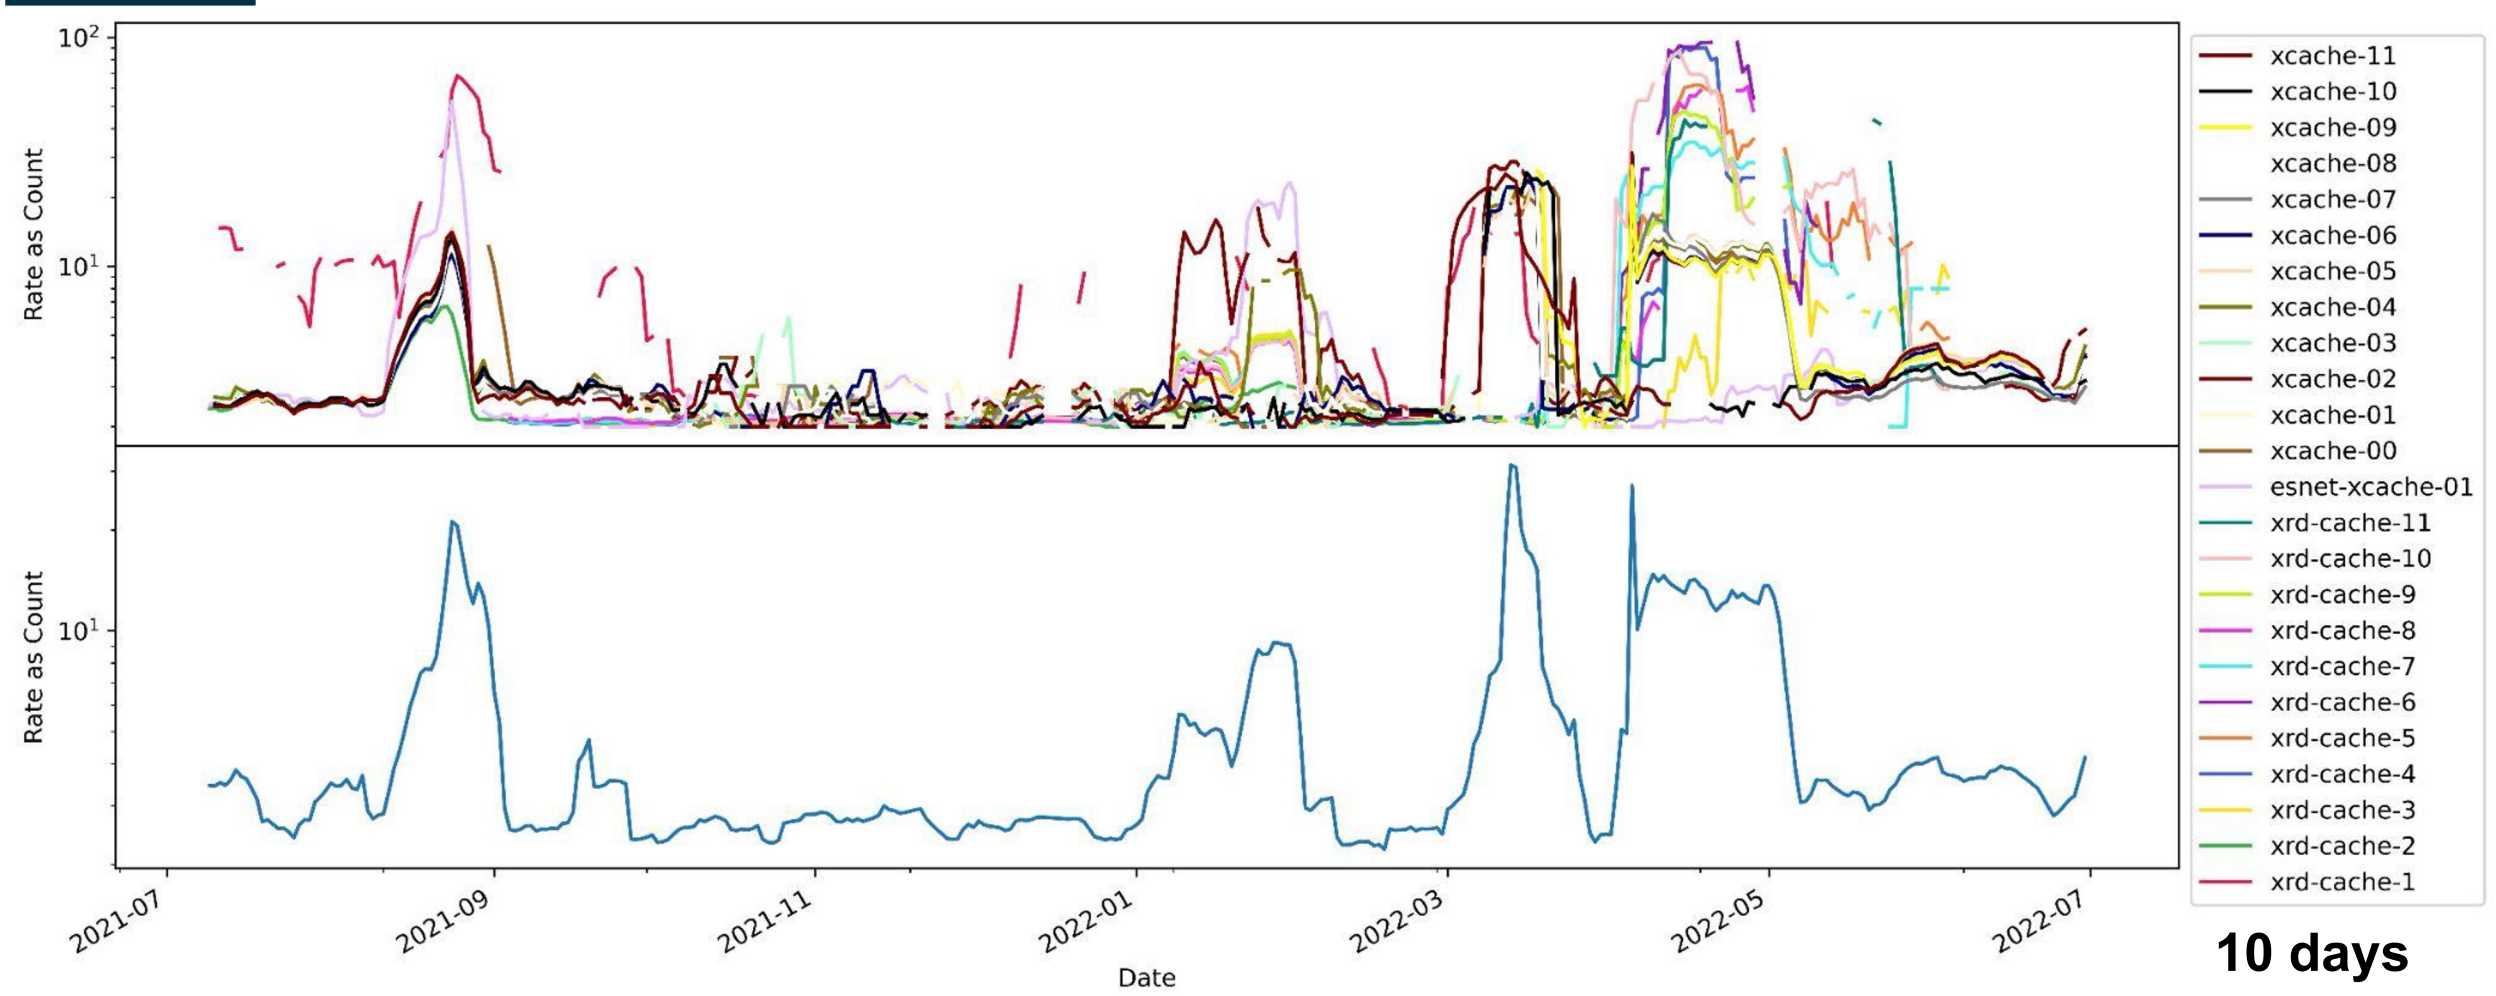

10 days

Daily data re-access rates with moving average 14 days (log scale)

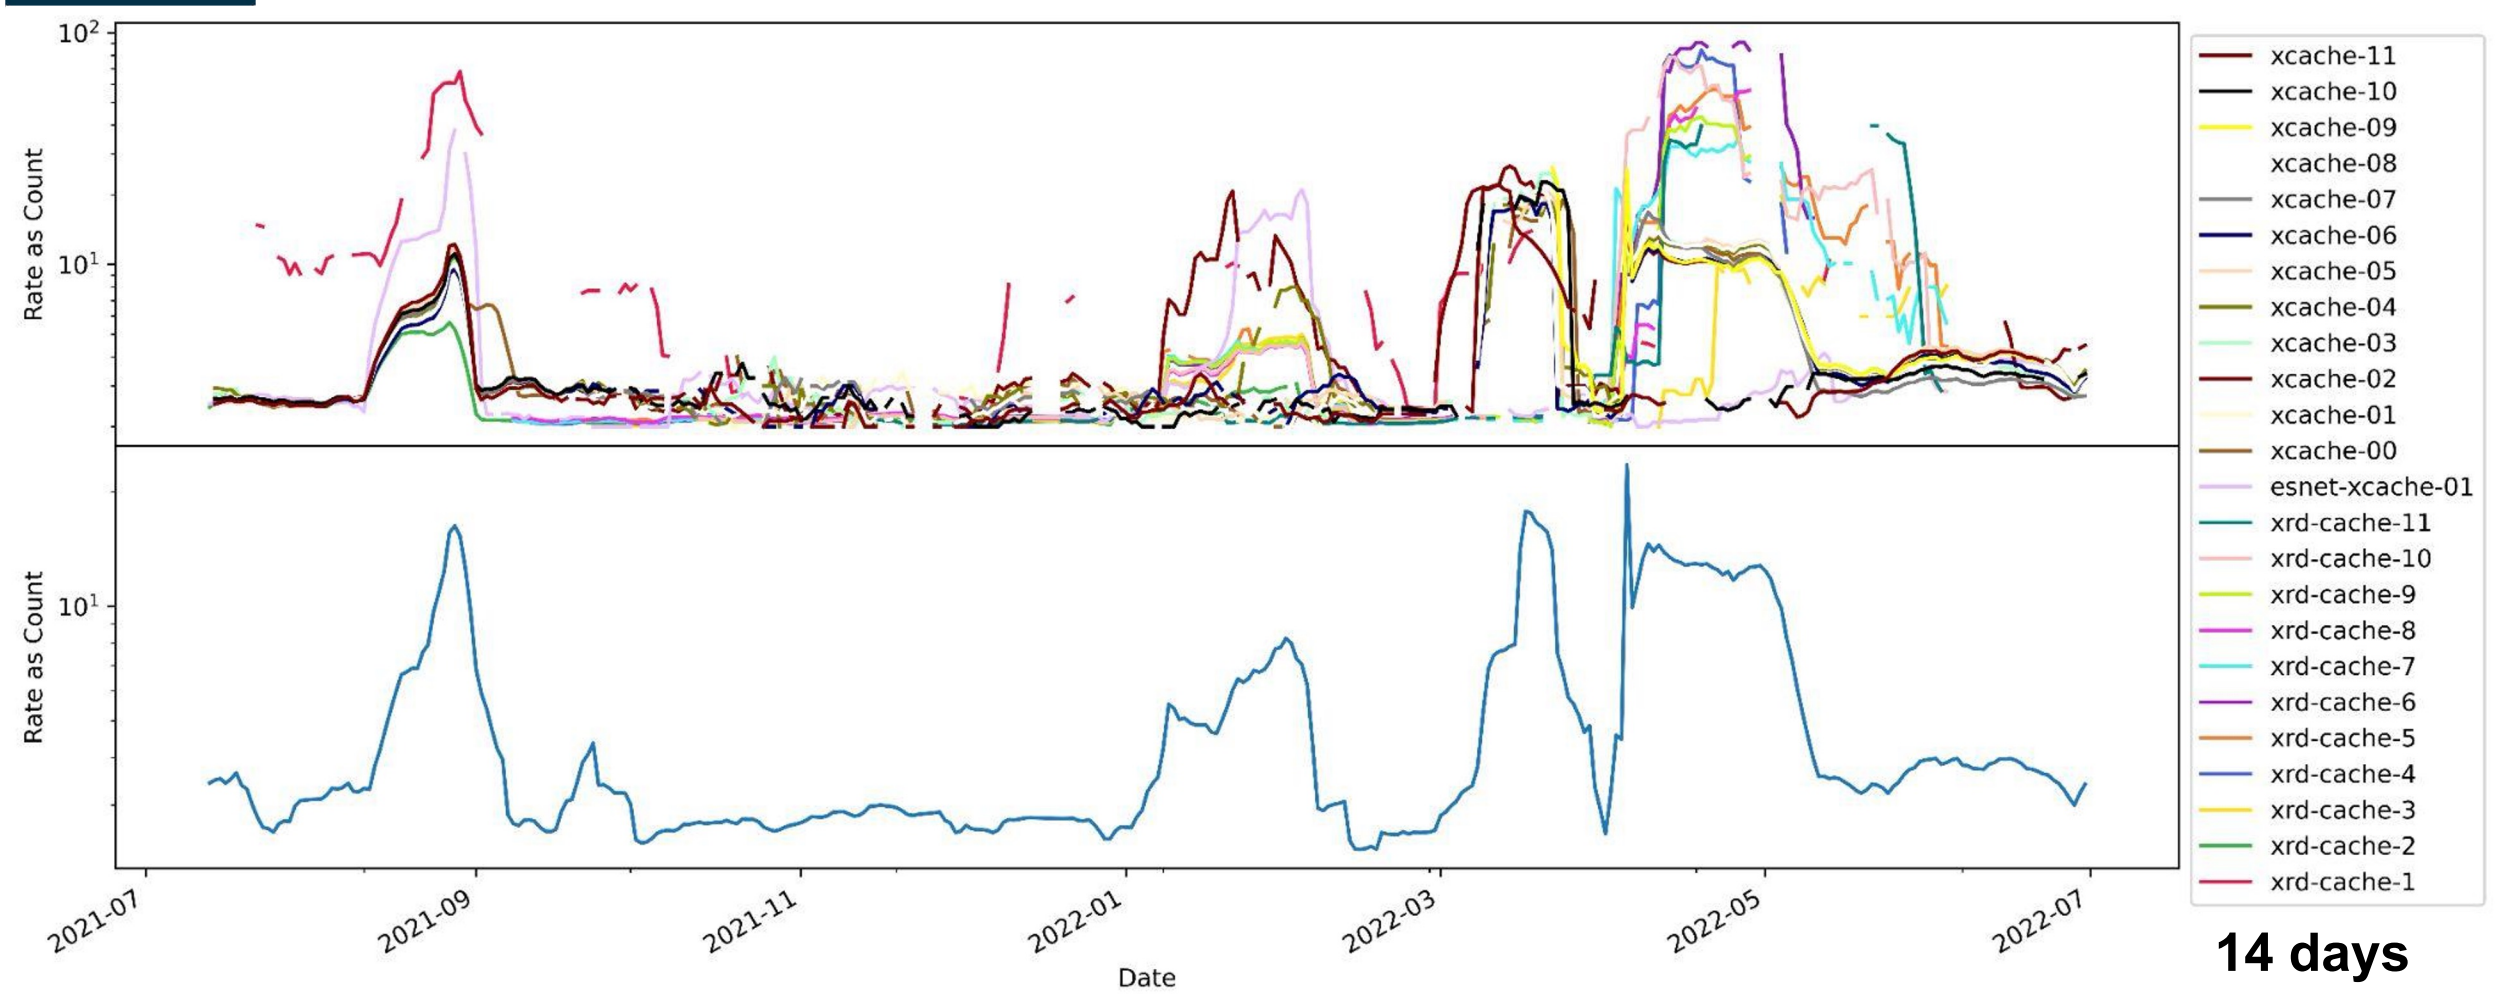

14 days

Daily data re-access rates with moving average 21 days (log scale)

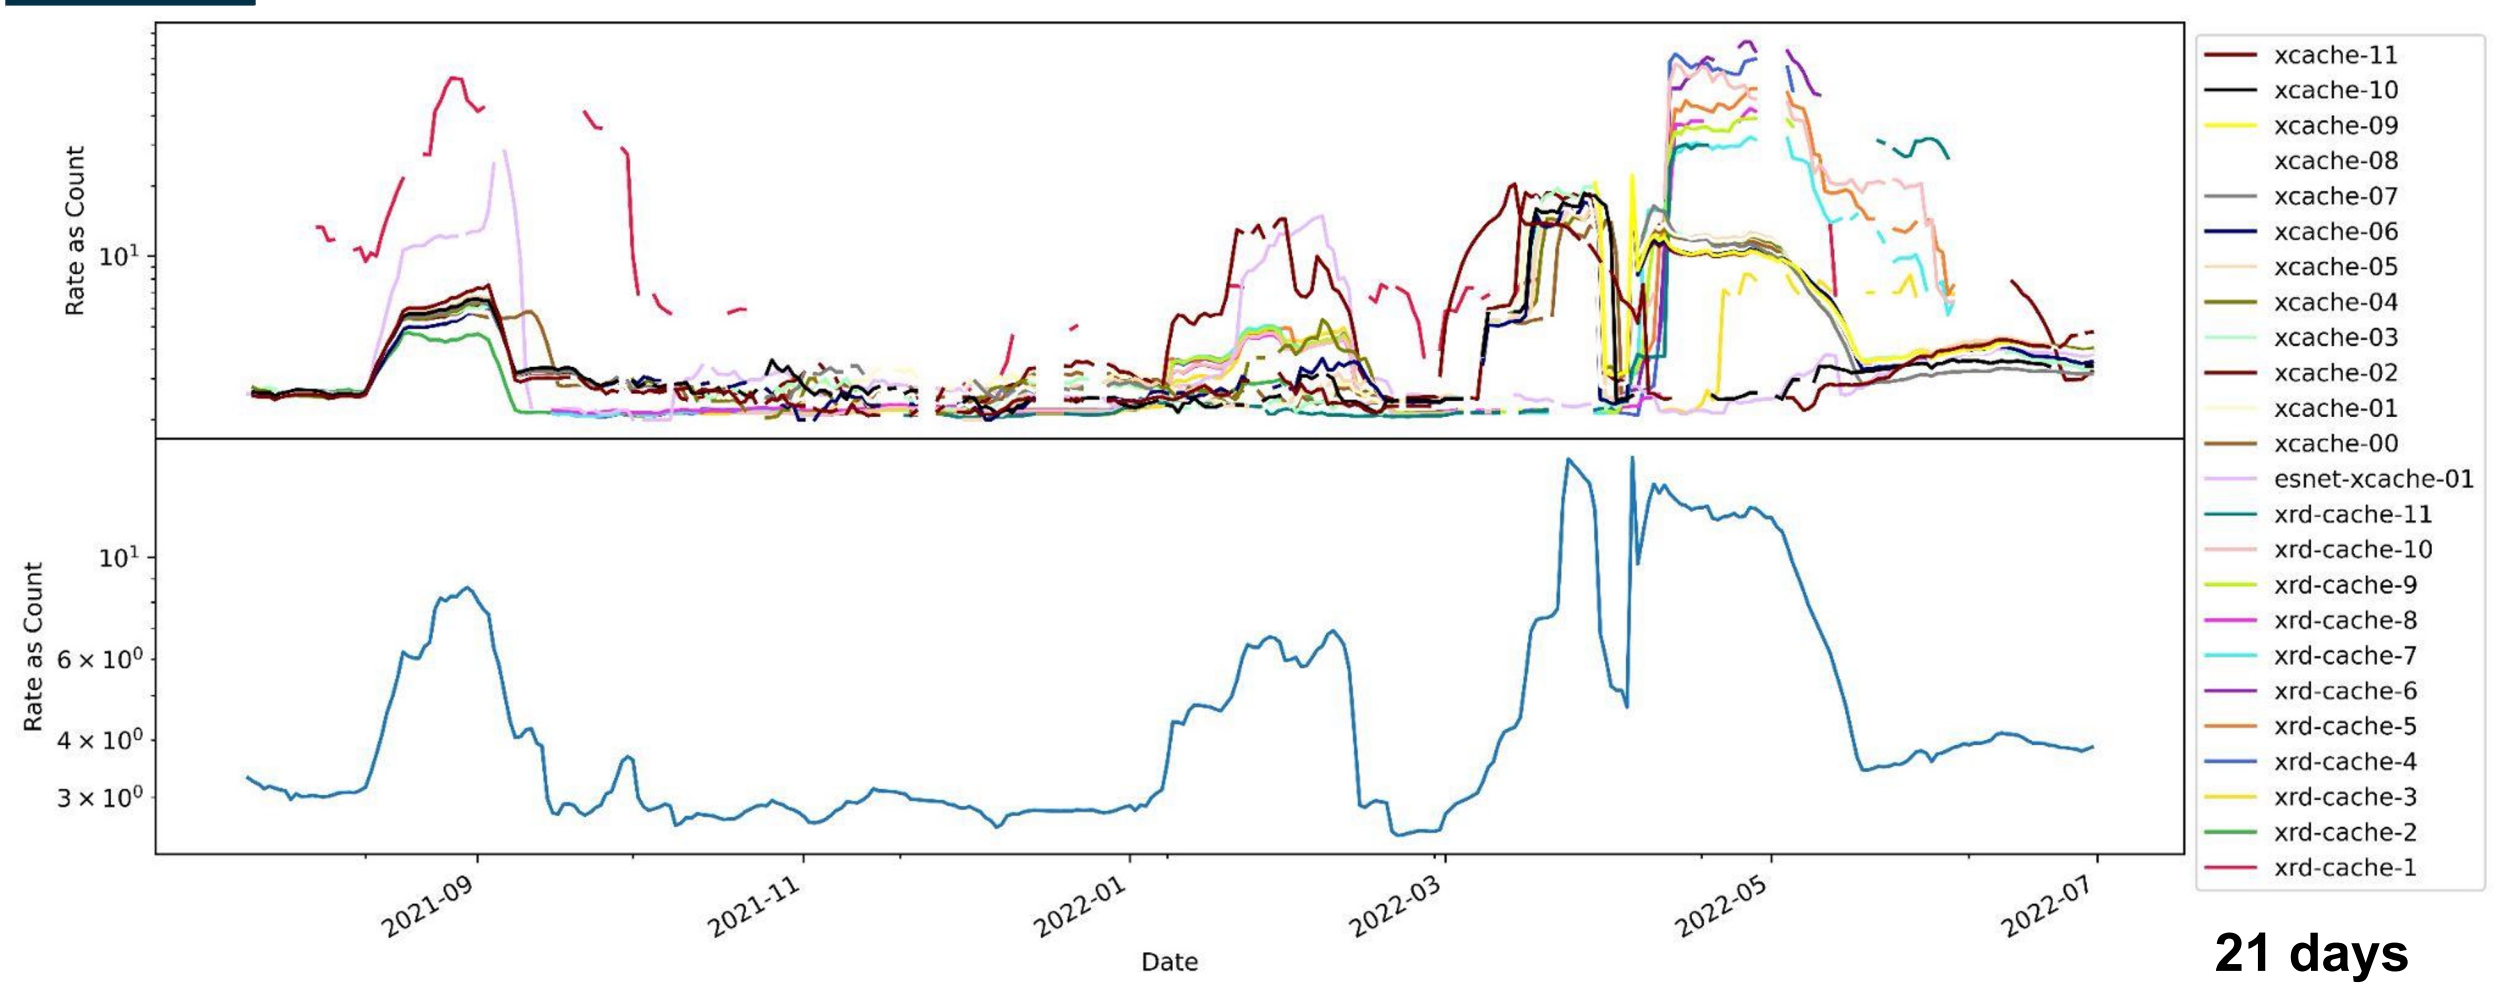

21 days

Daily data re-access rates with moving average 28 days (log scale)

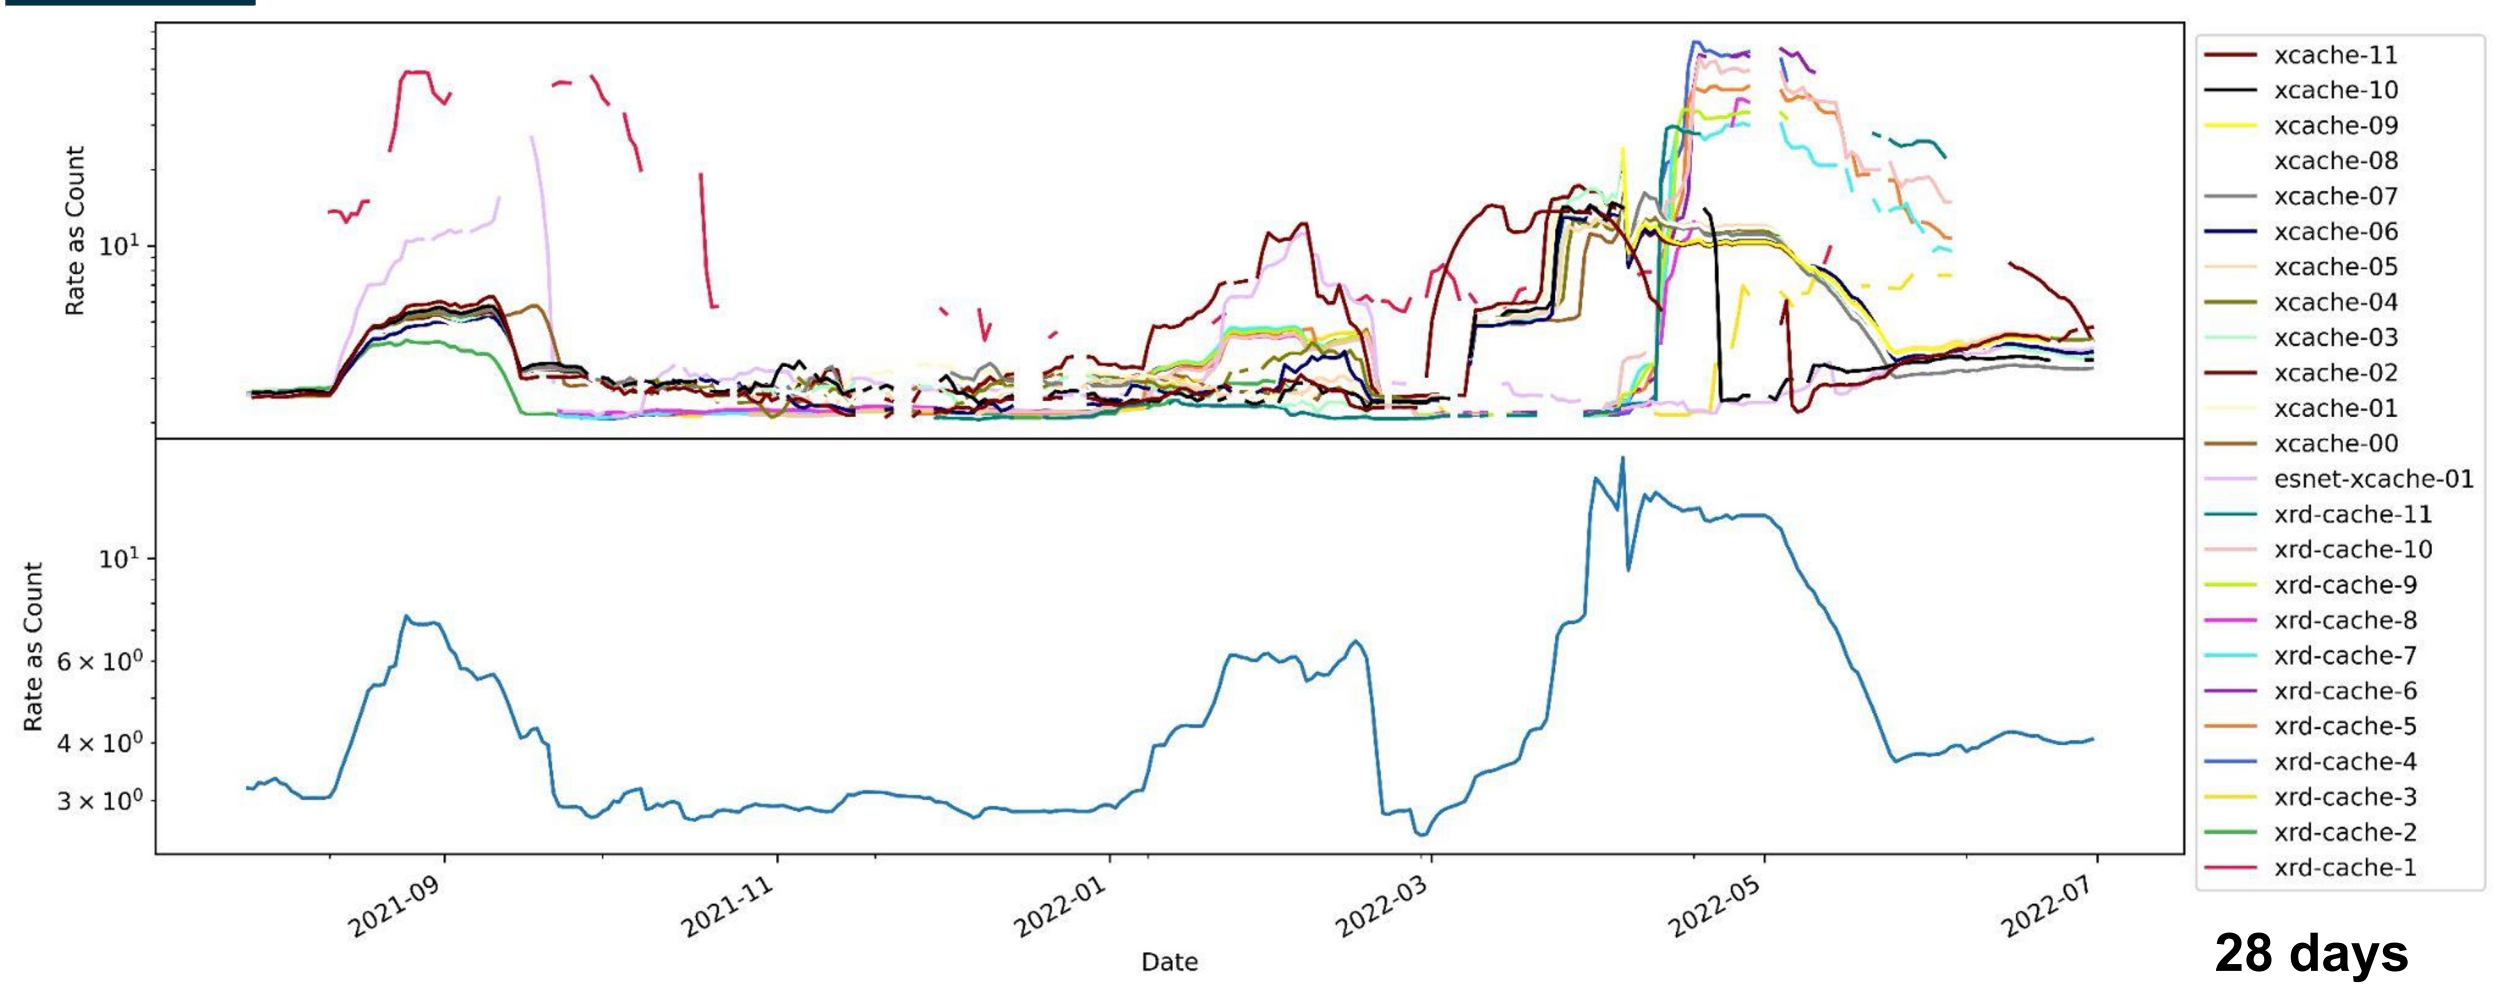

28 days

Daily data re-access volume

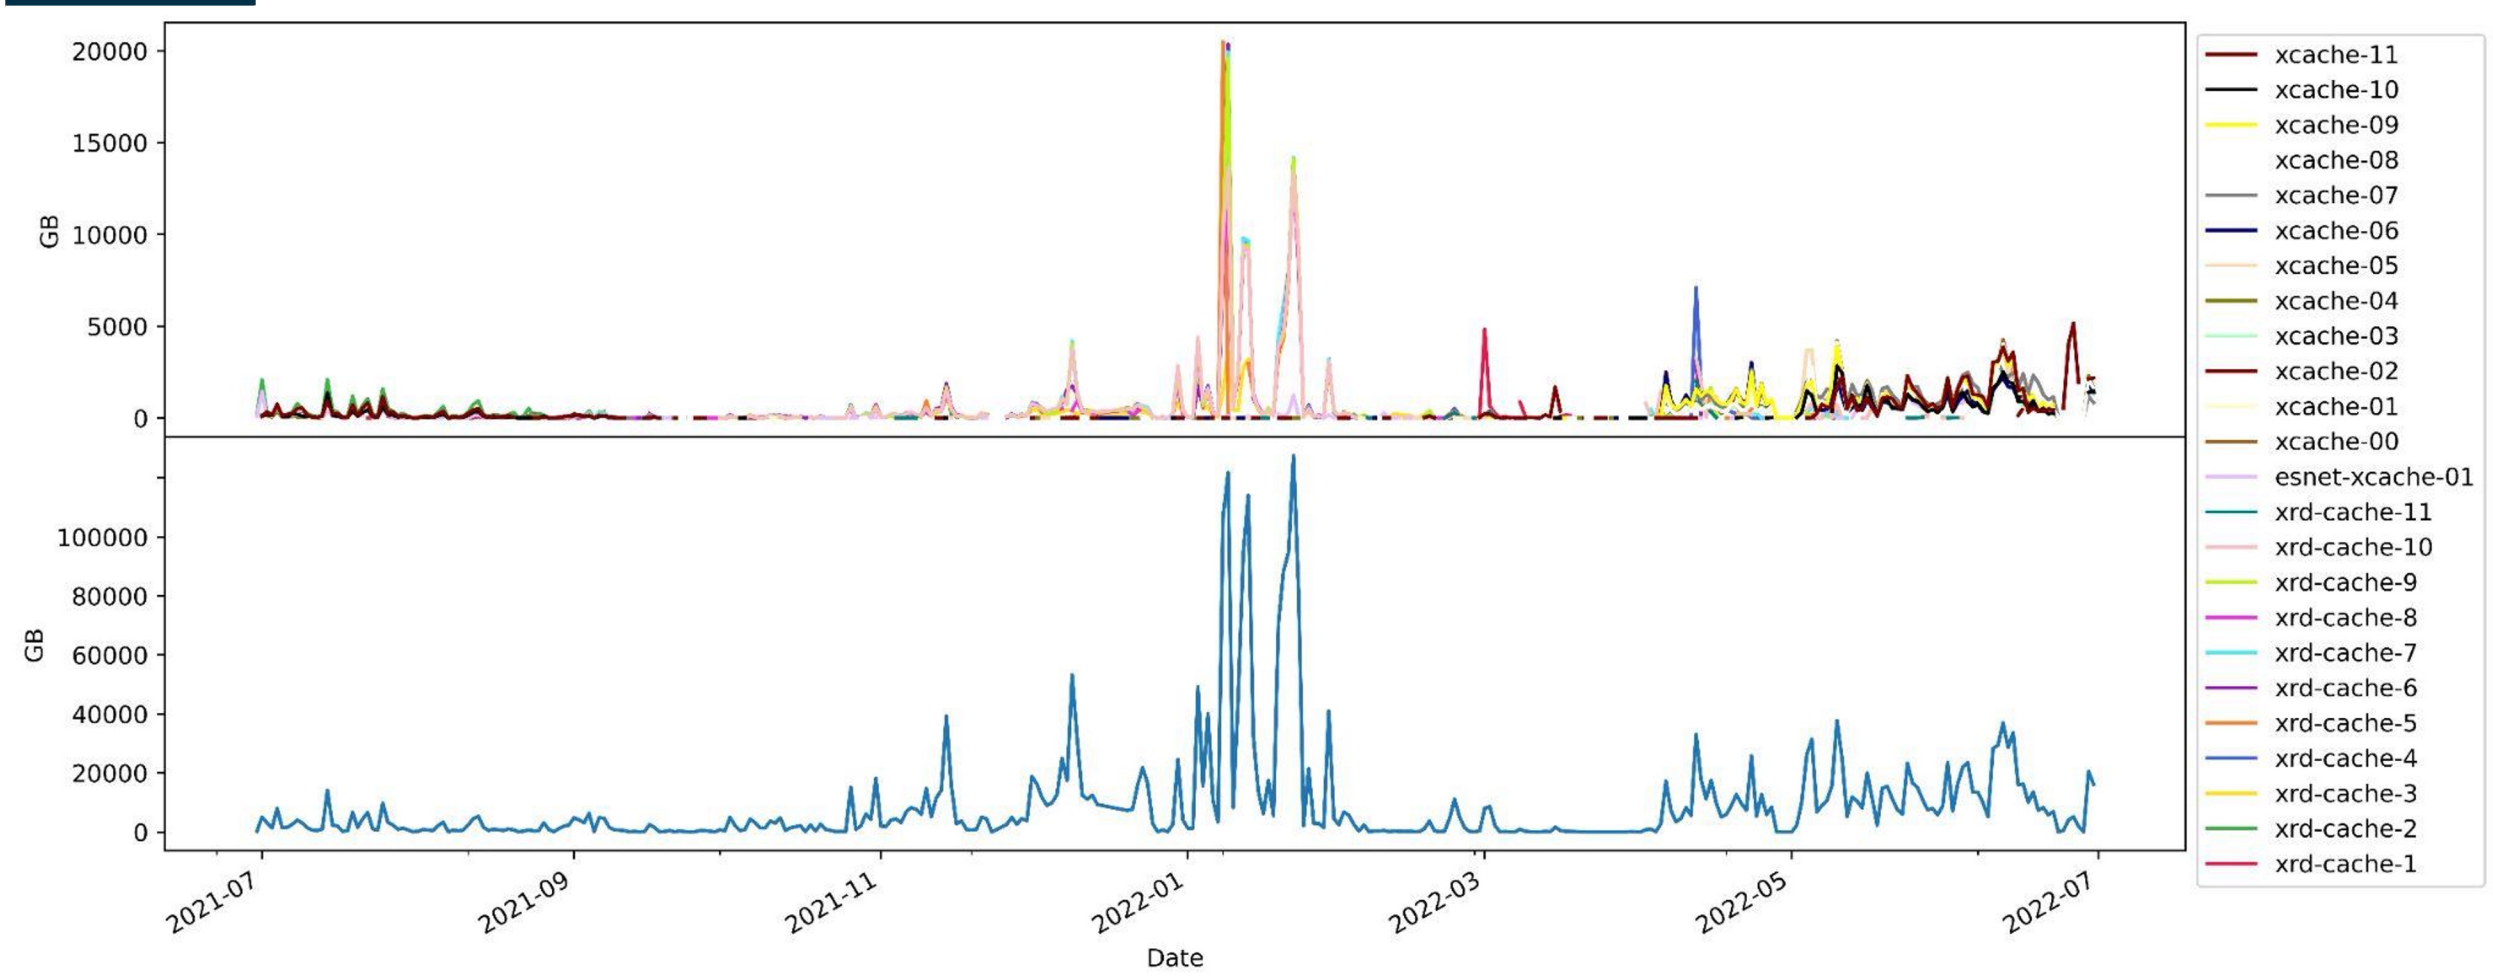

Daily data re-access volume with 2-day moving average

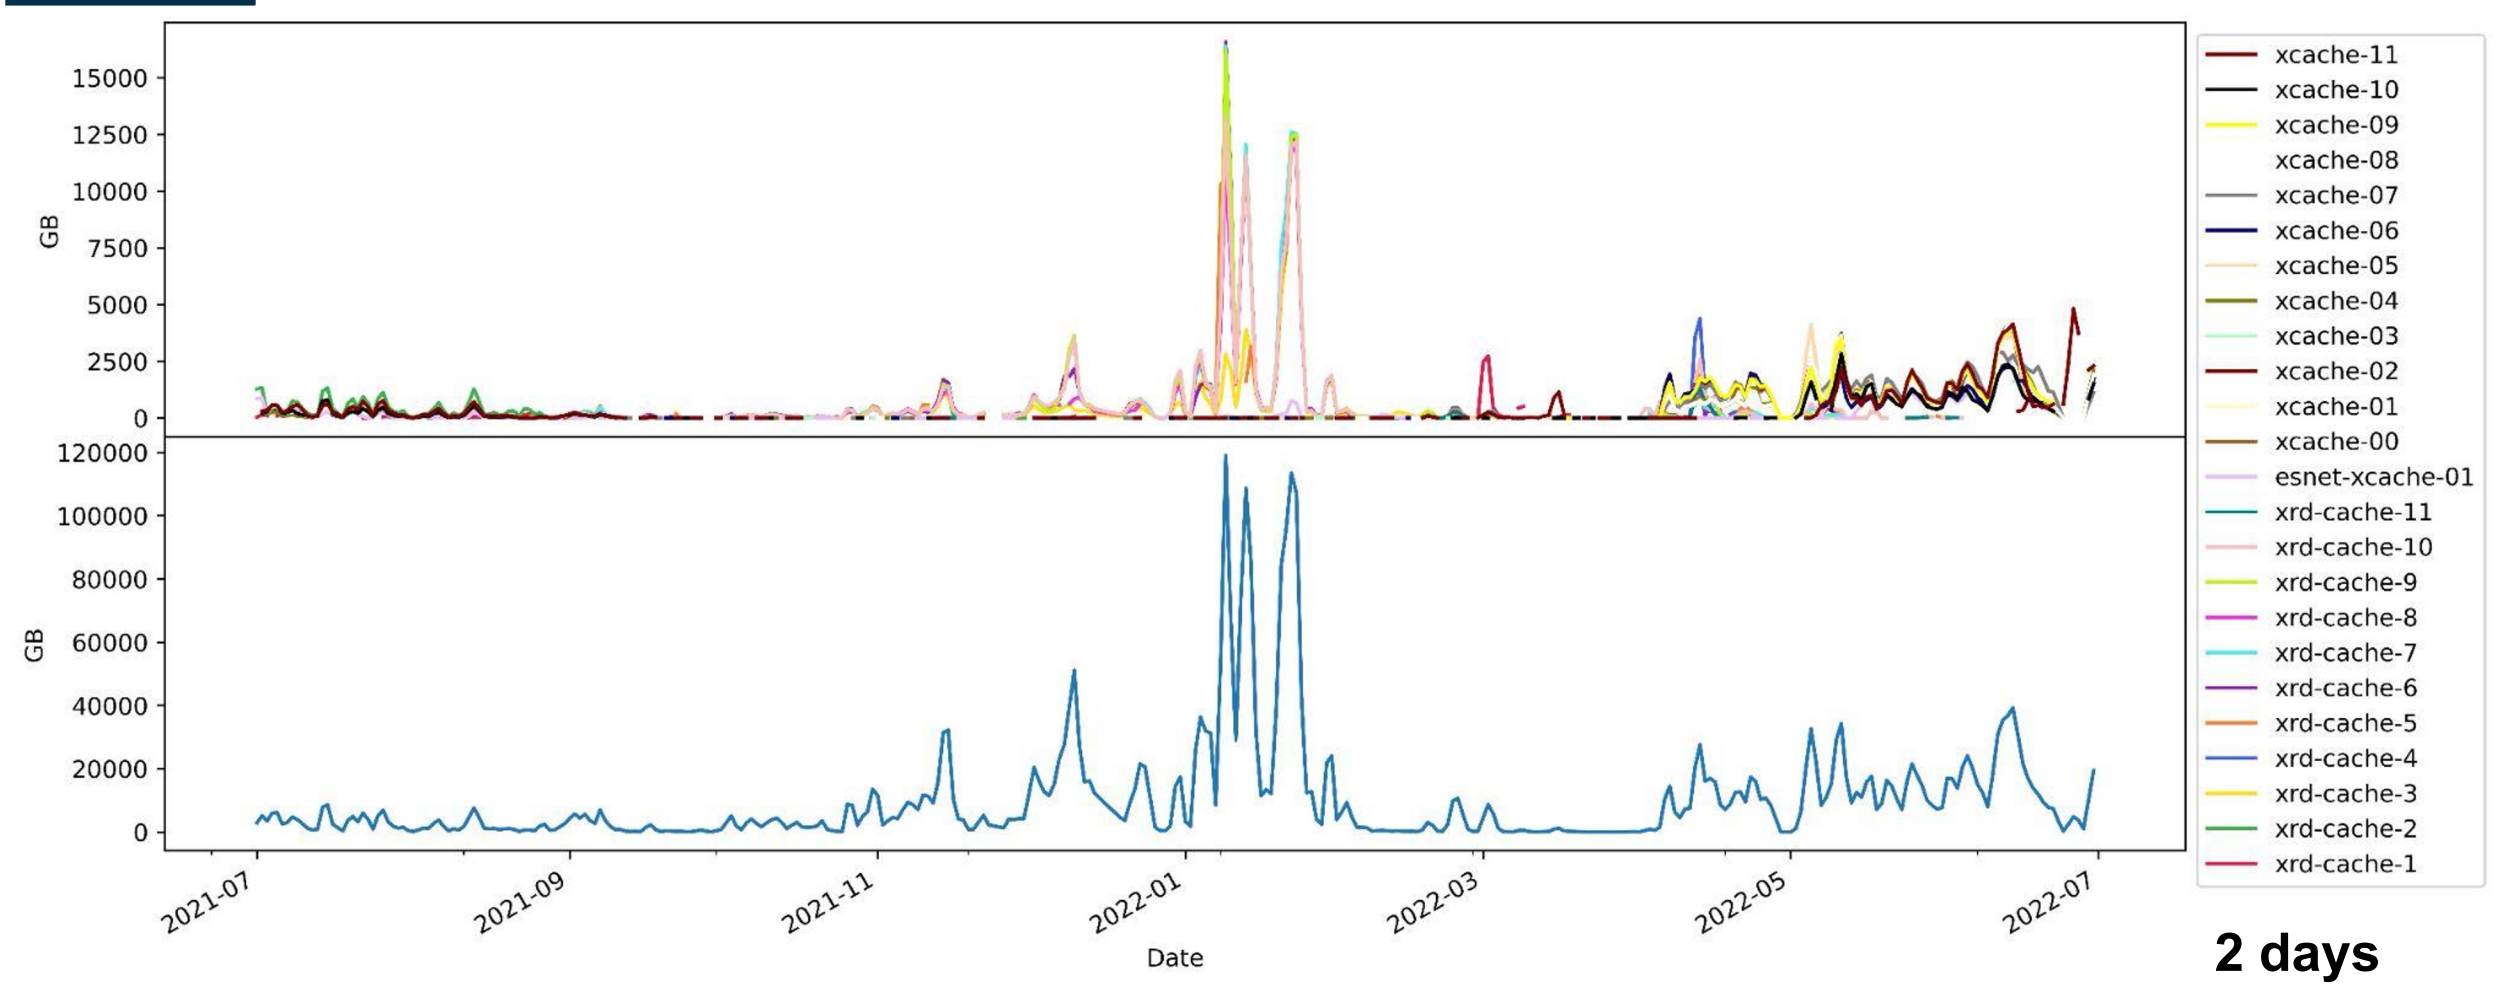

2 days

Daily data re-access volume with 3-day moving average

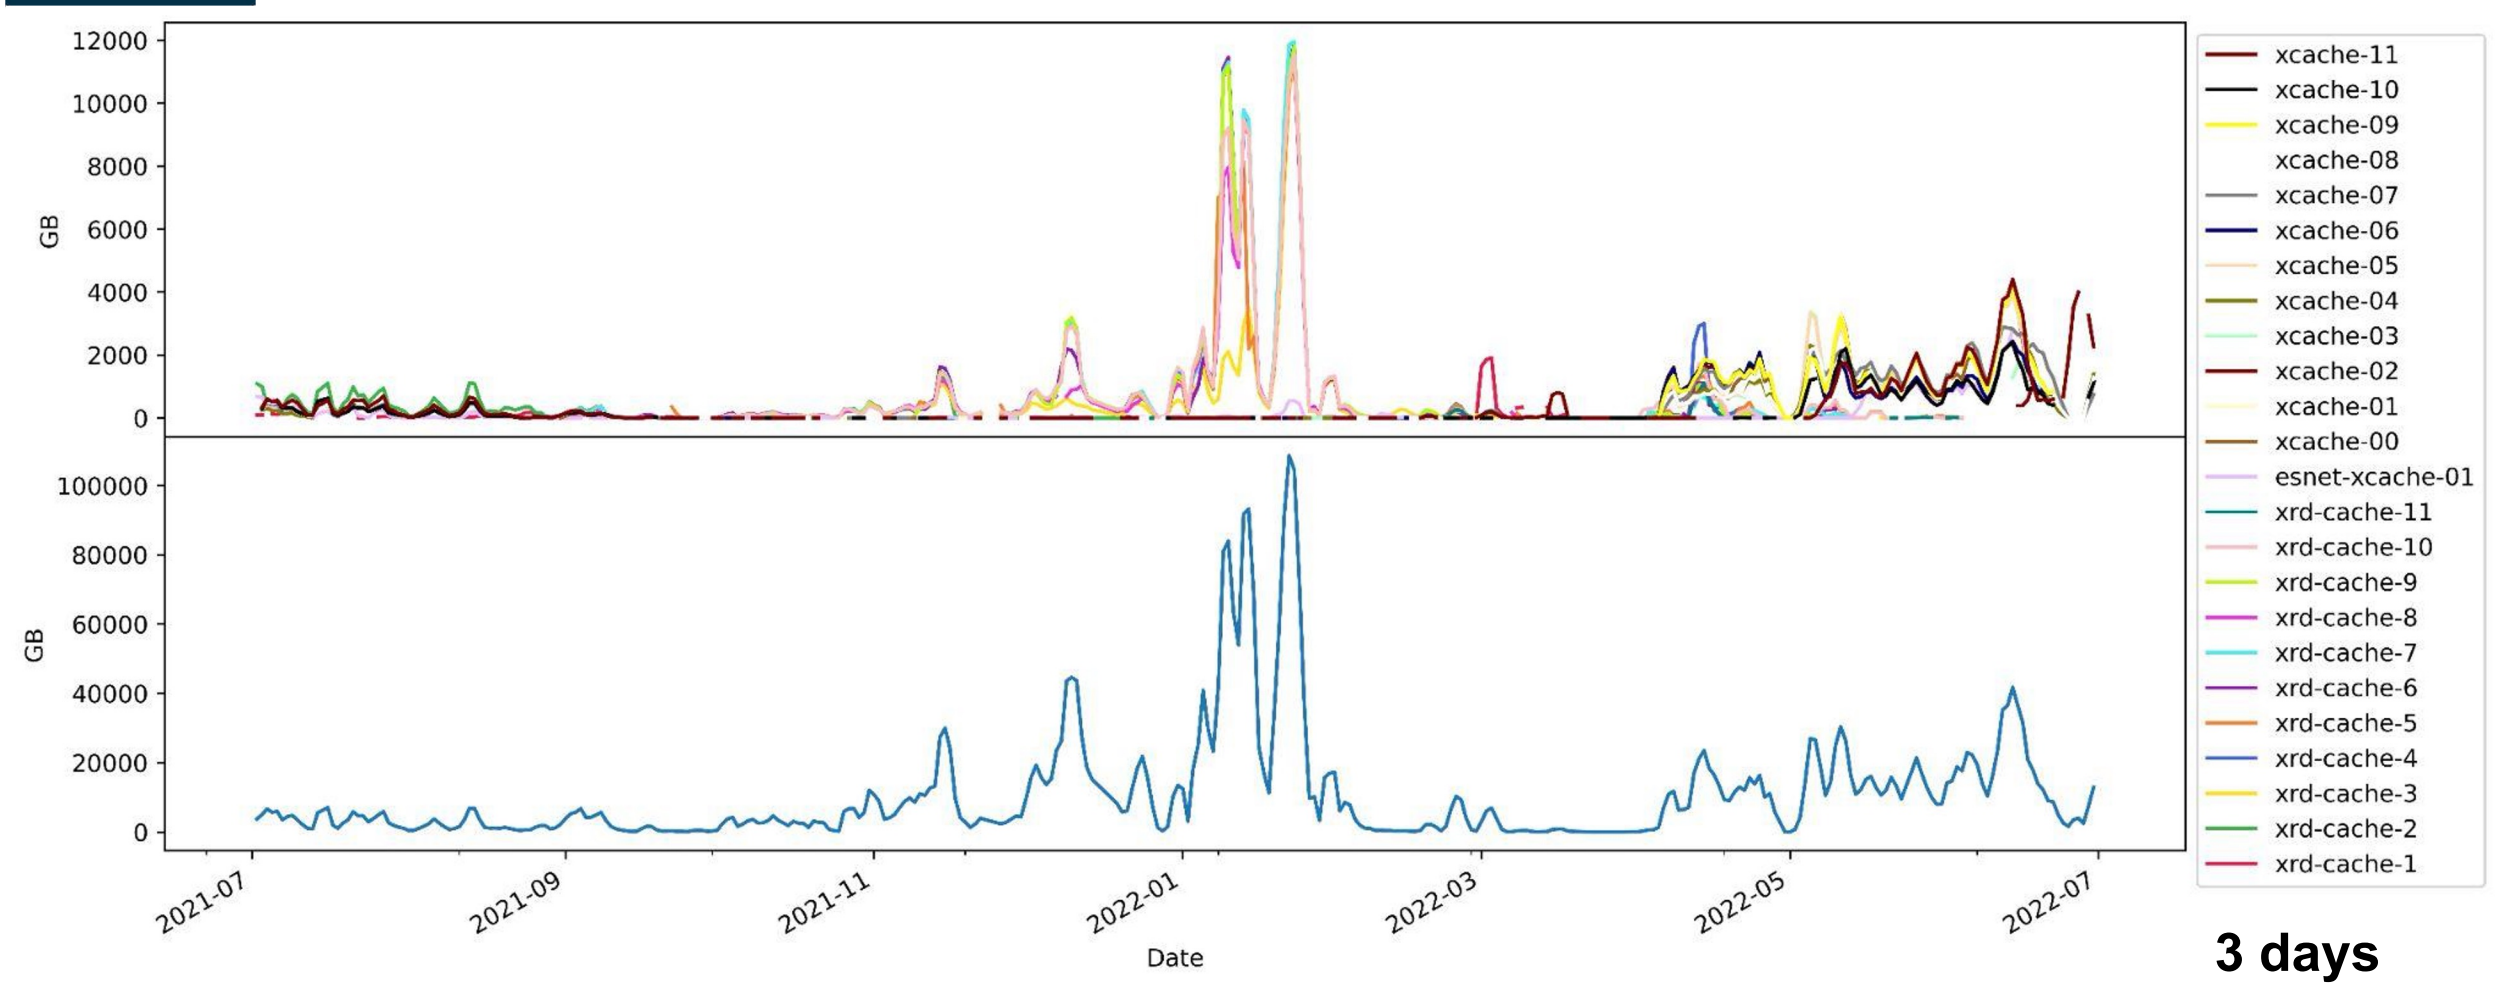

3 days

Daily data re-access volume with 5-day moving average

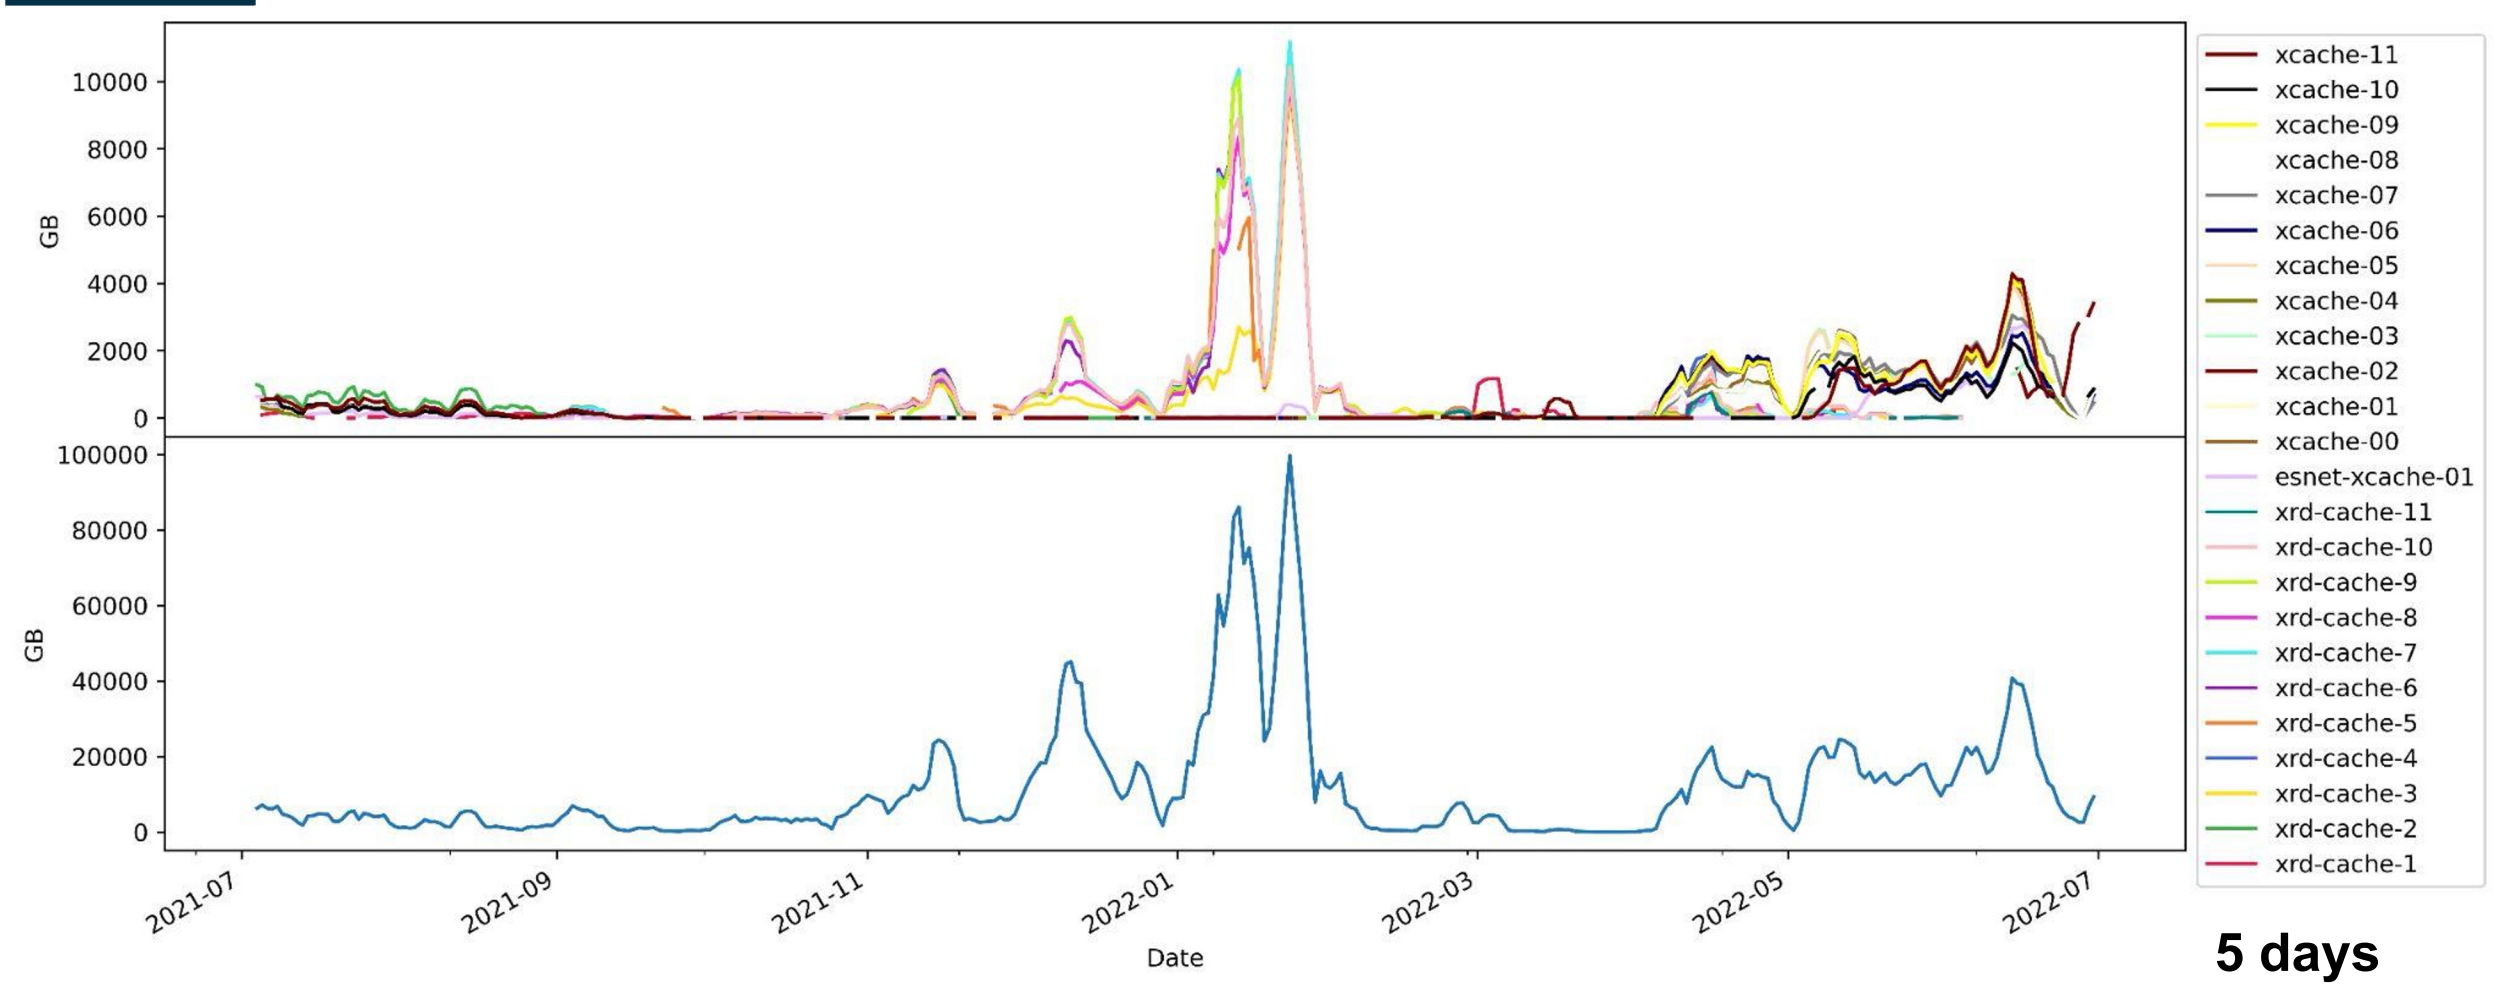

5 days

Daily data re-access volume with 7-day moving average

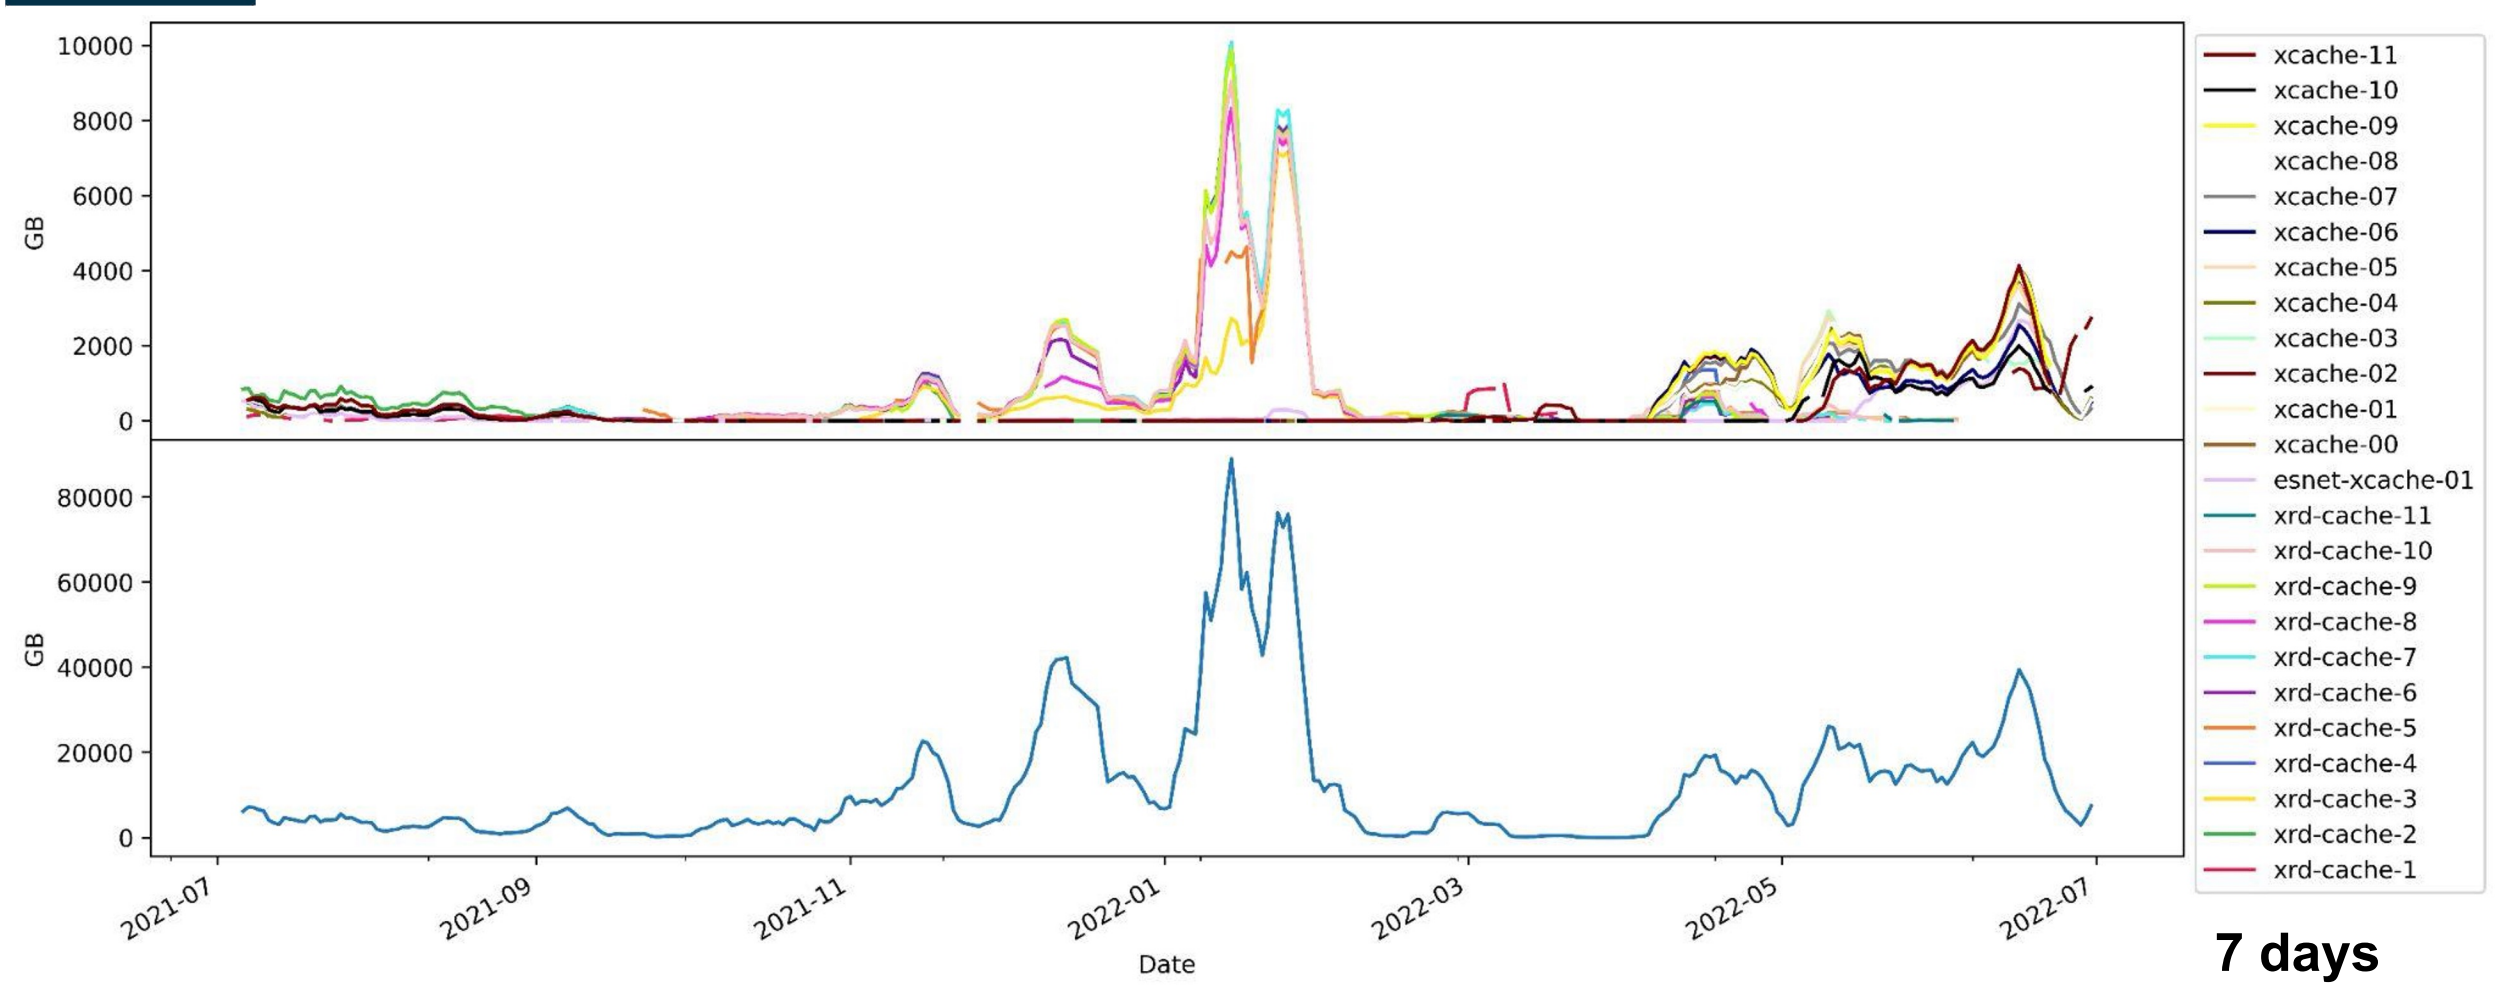

7 days

Daily data re-access volume with 10-day moving average

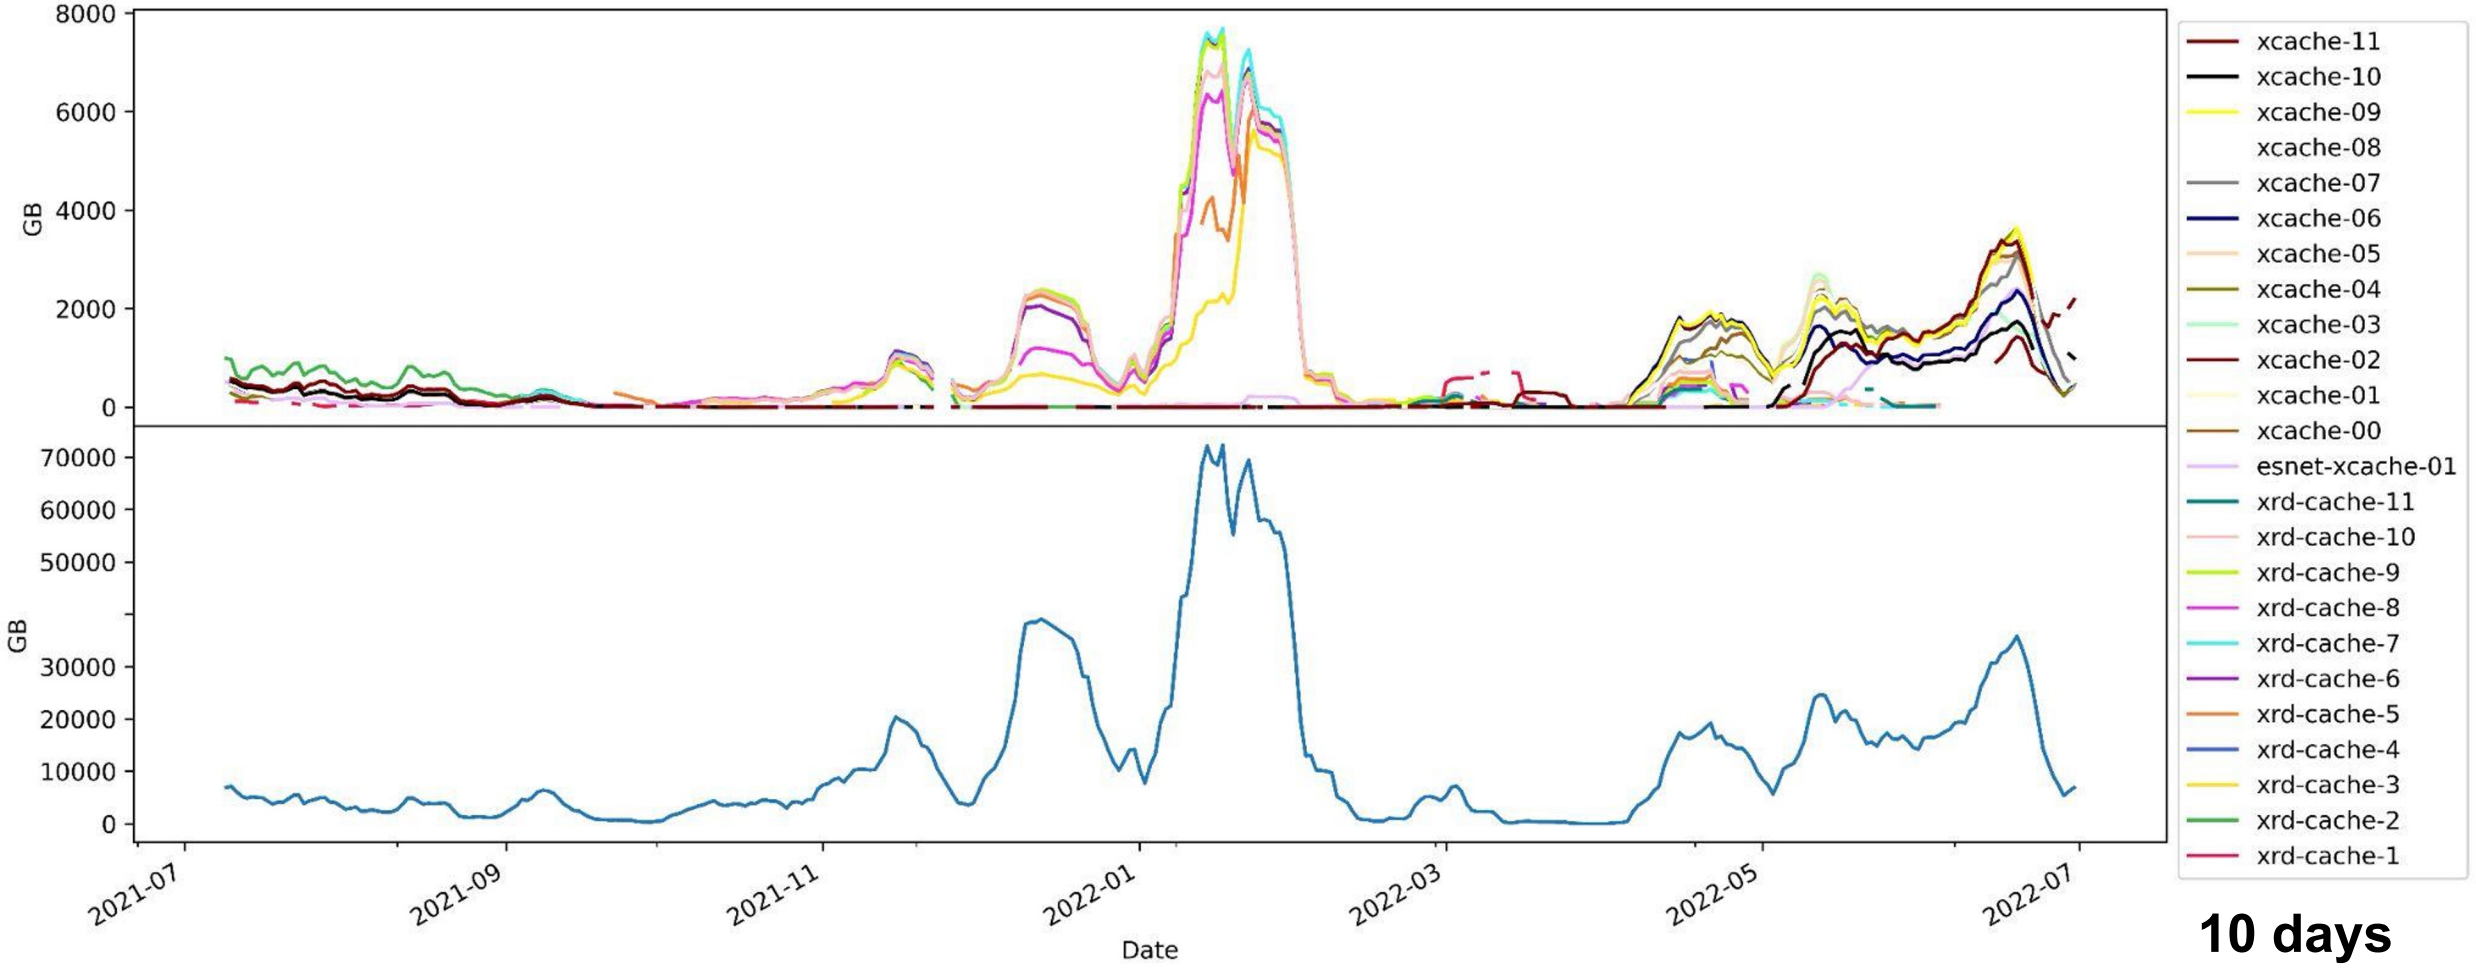

10 days

Daily data re-access volume with 14-day moving average

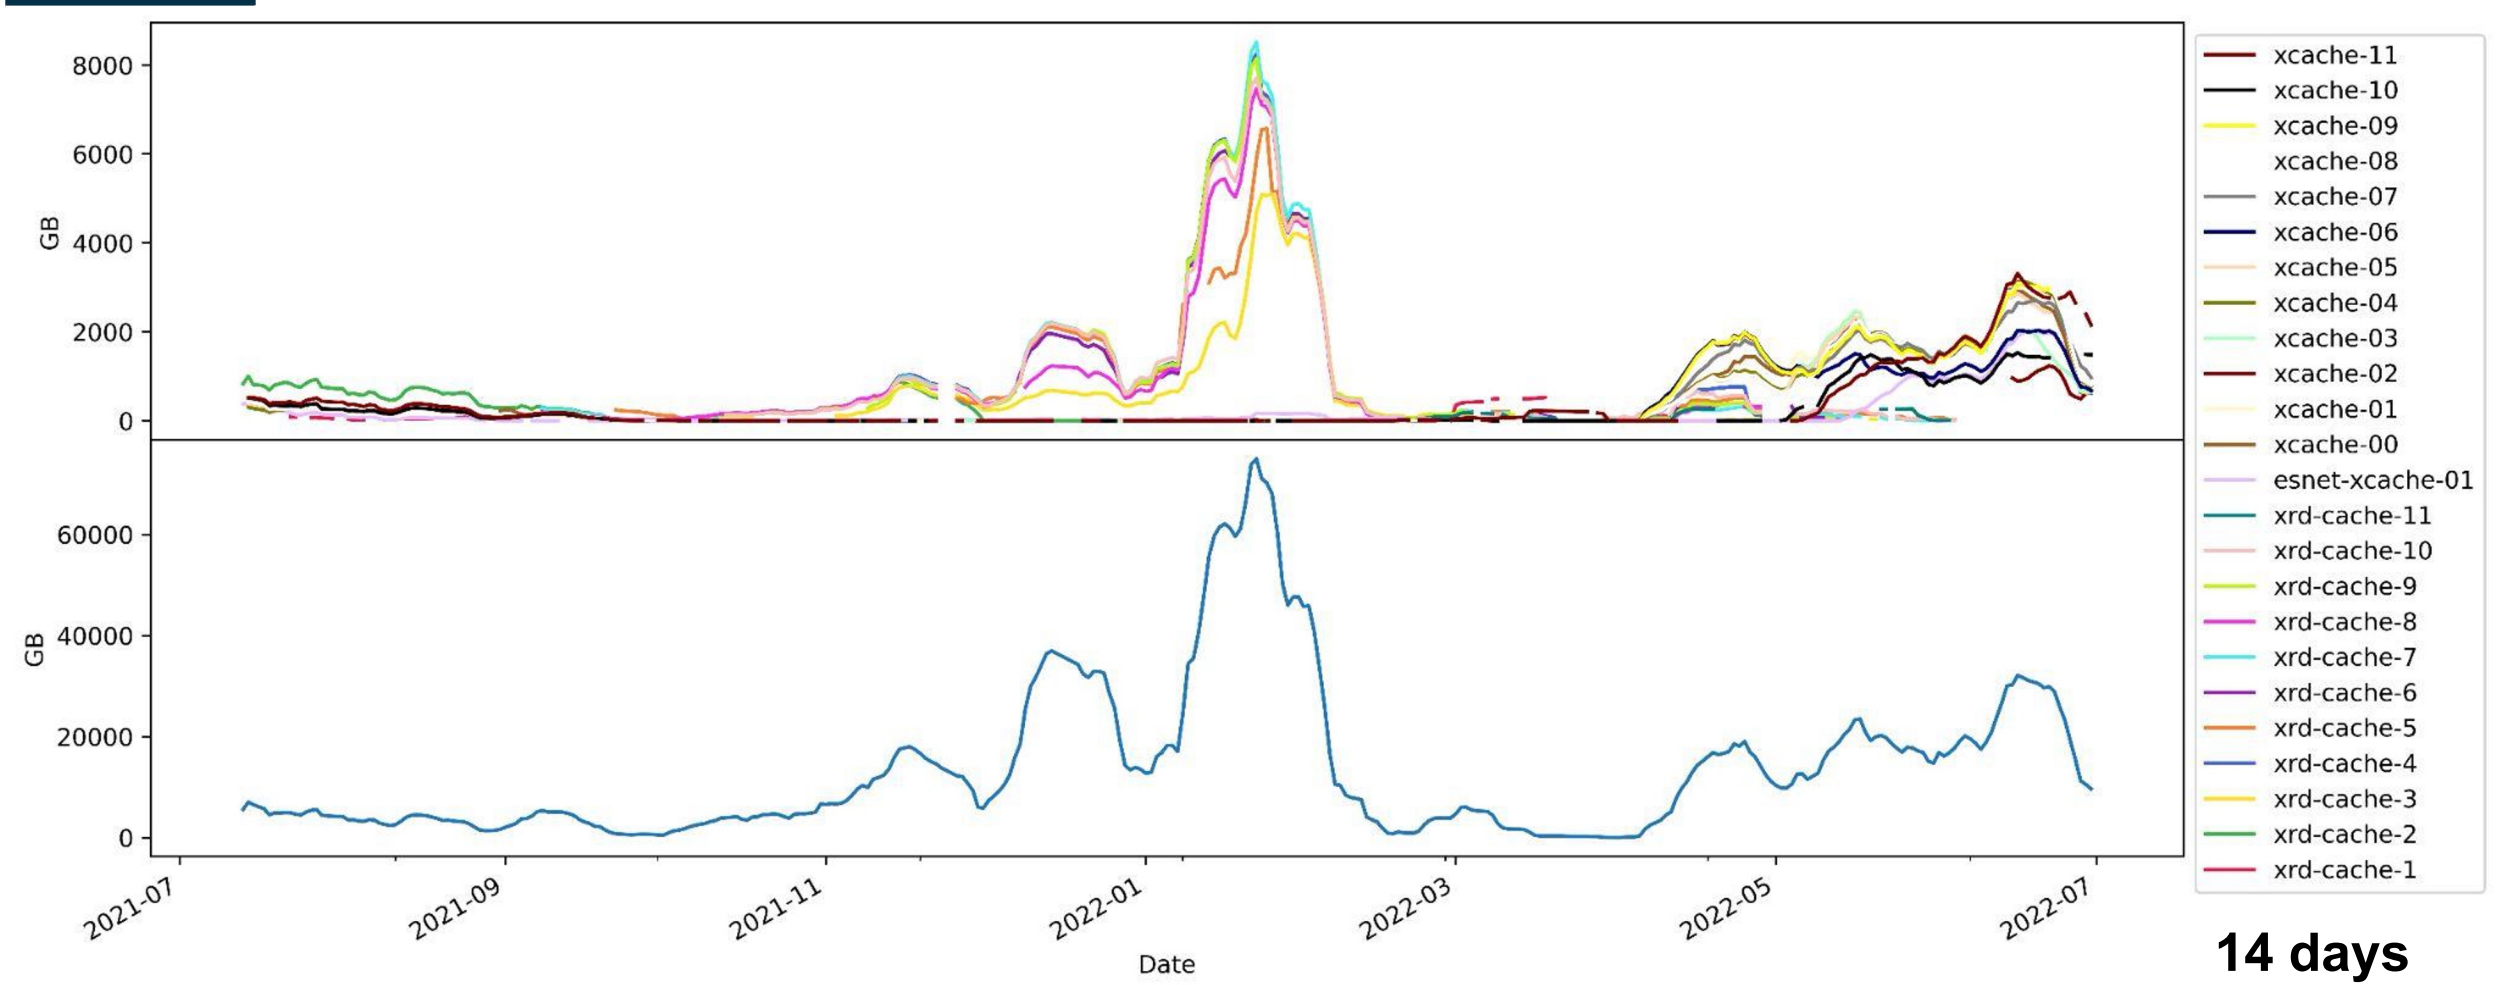

14 days

Daily data re-access volume with 21-day moving average

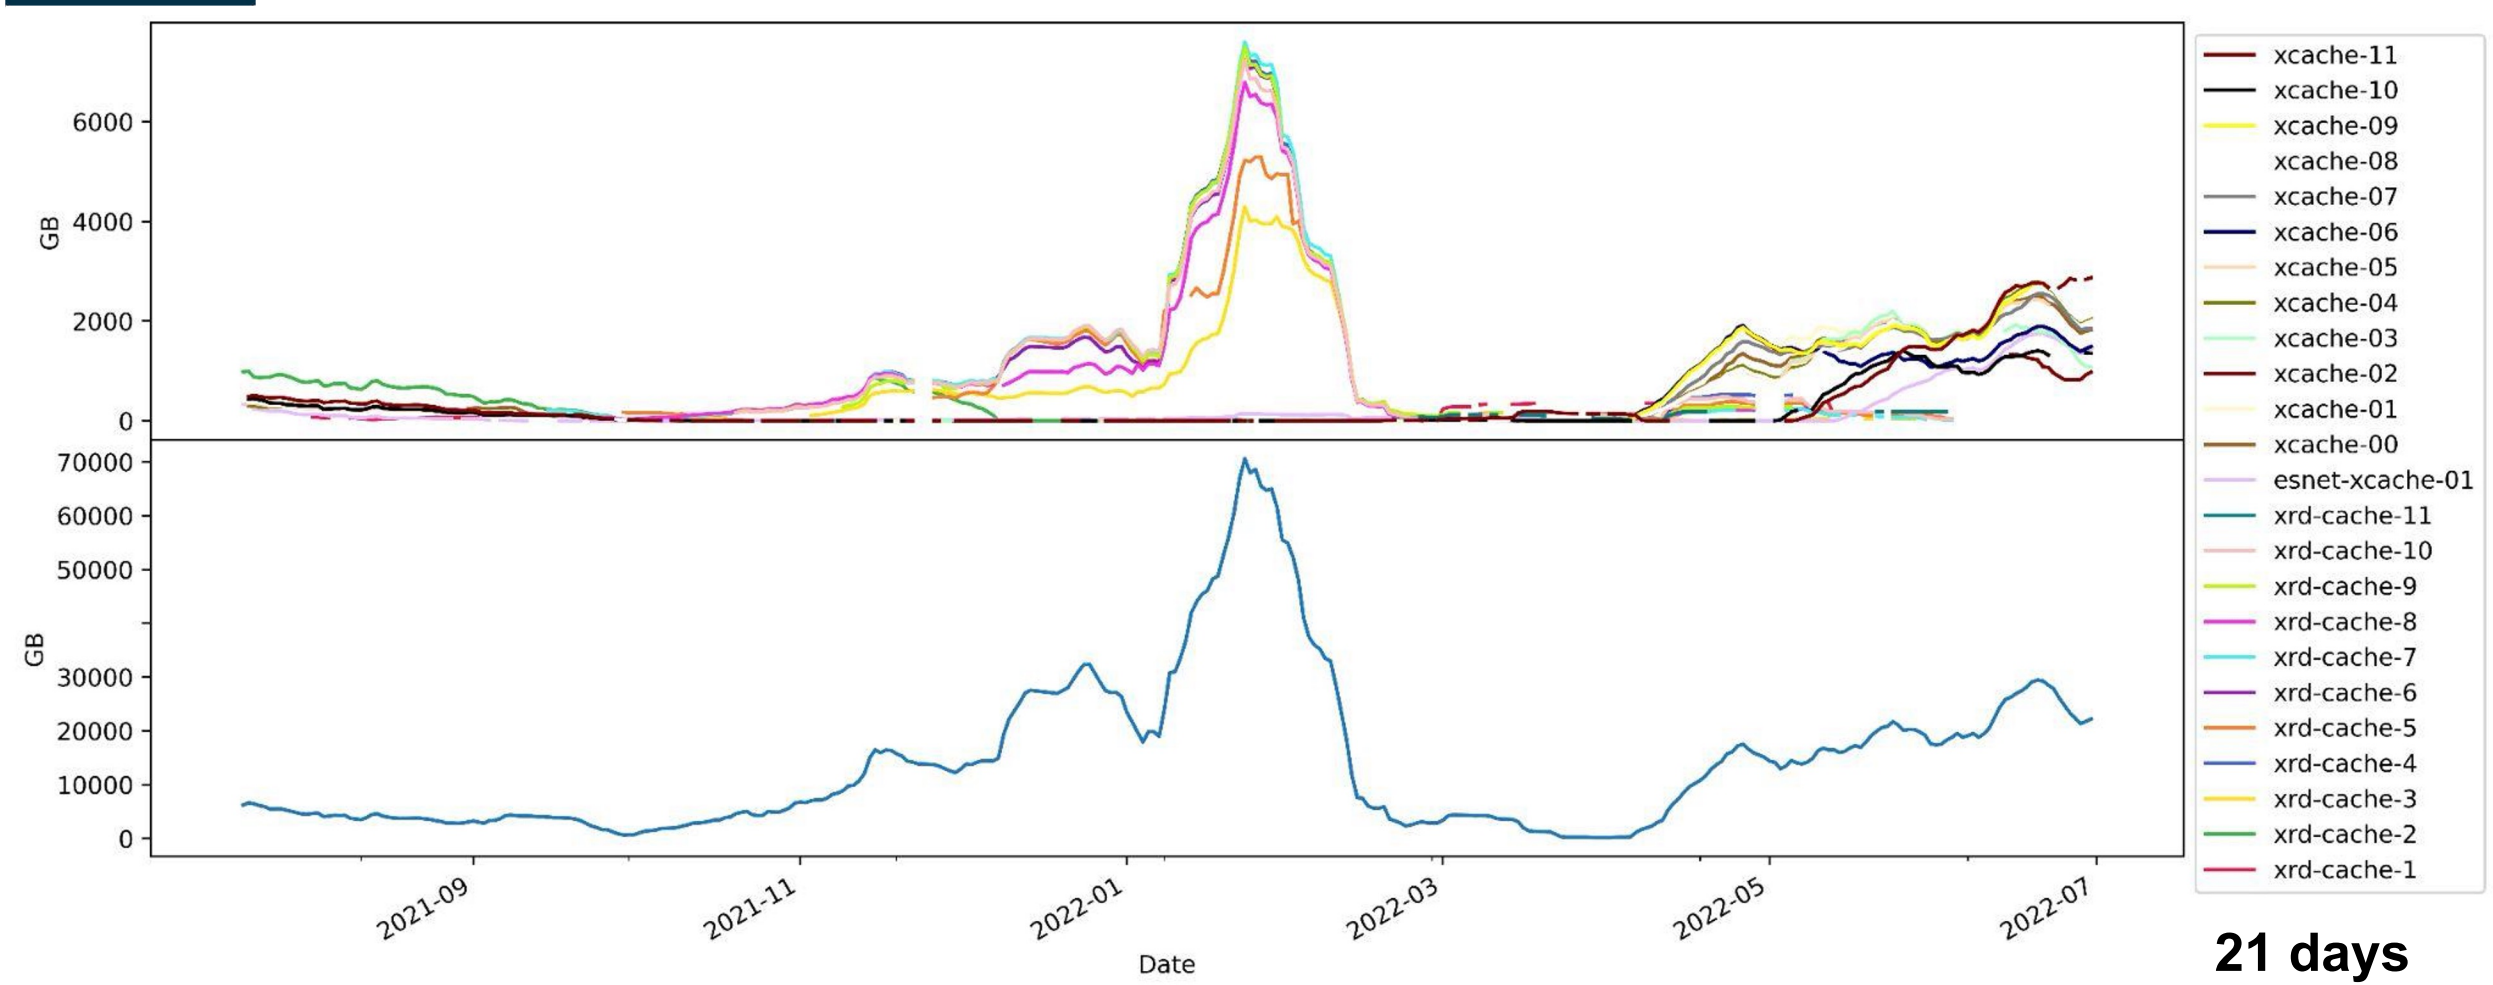

21 days

Daily data re-access volume with 28-day moving average

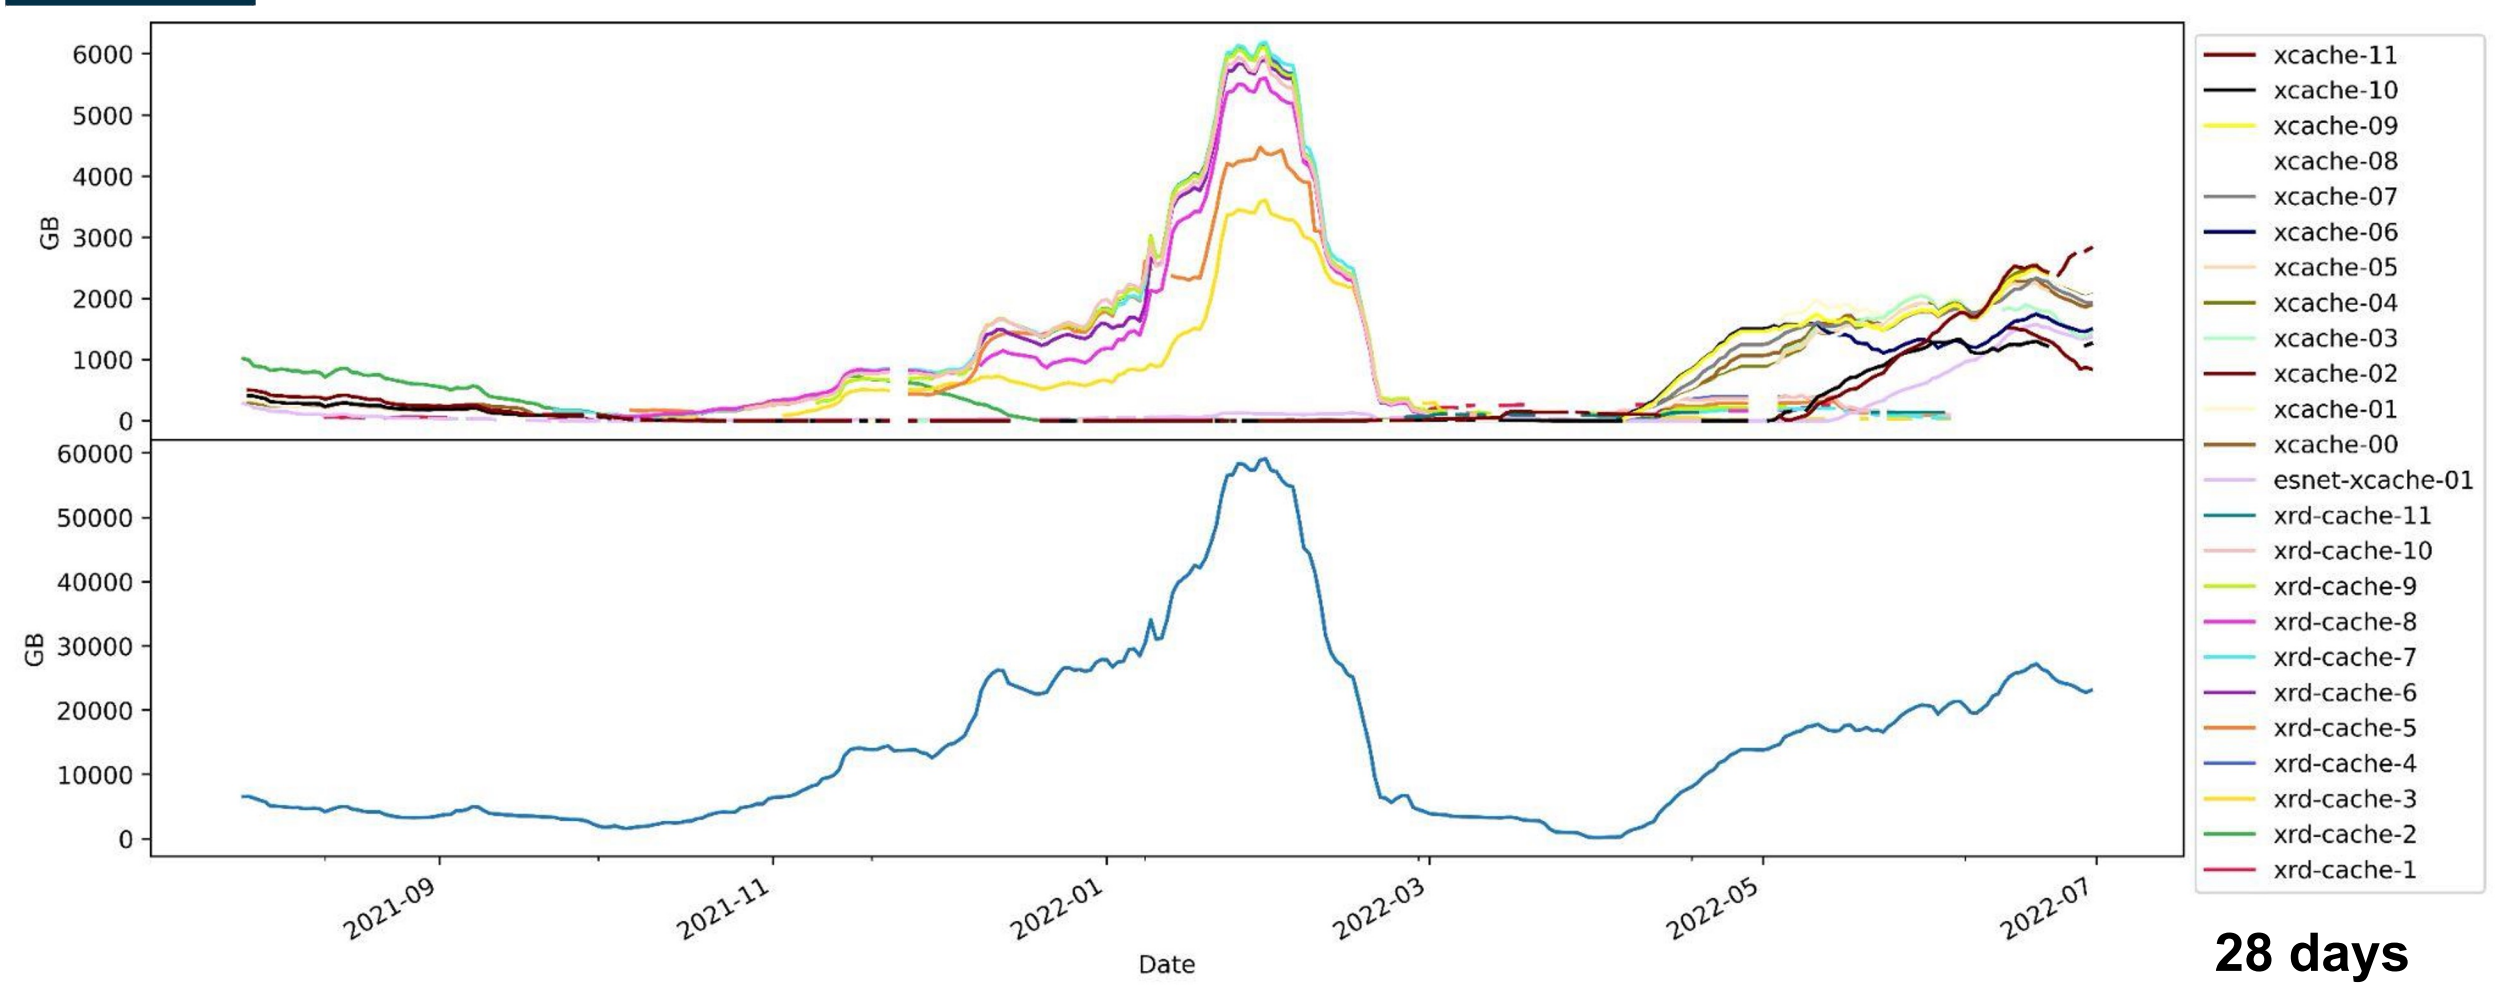

28 days

Data access throughput performance

Daily average data access throughput performance for cache hit per access

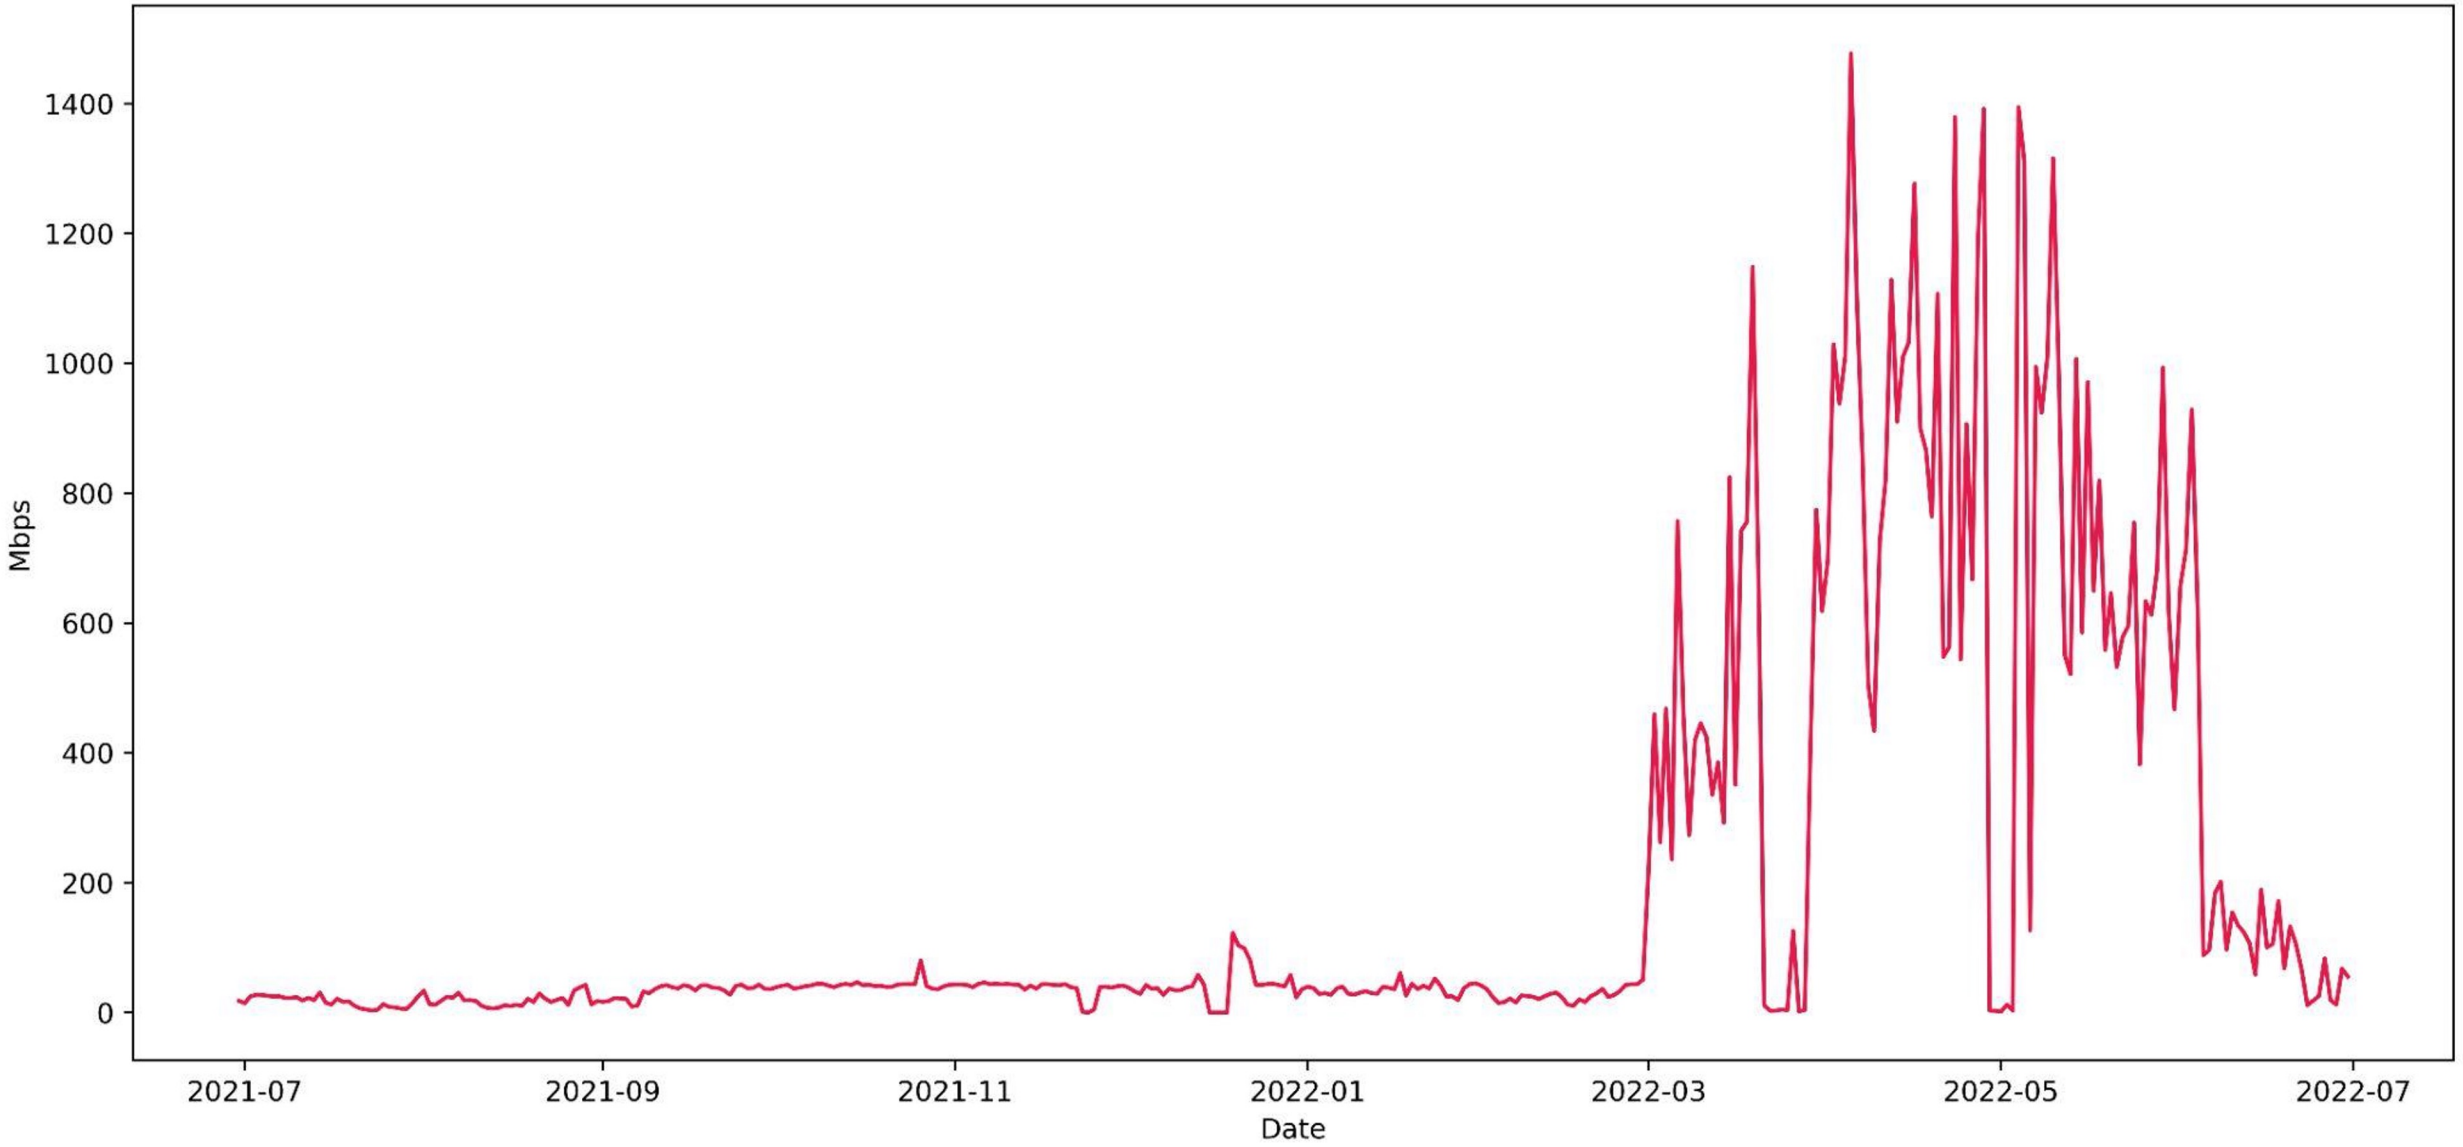

Daily cumulative data access throughput performance for cache hits

Daily average data throughput performance for cache miss per access

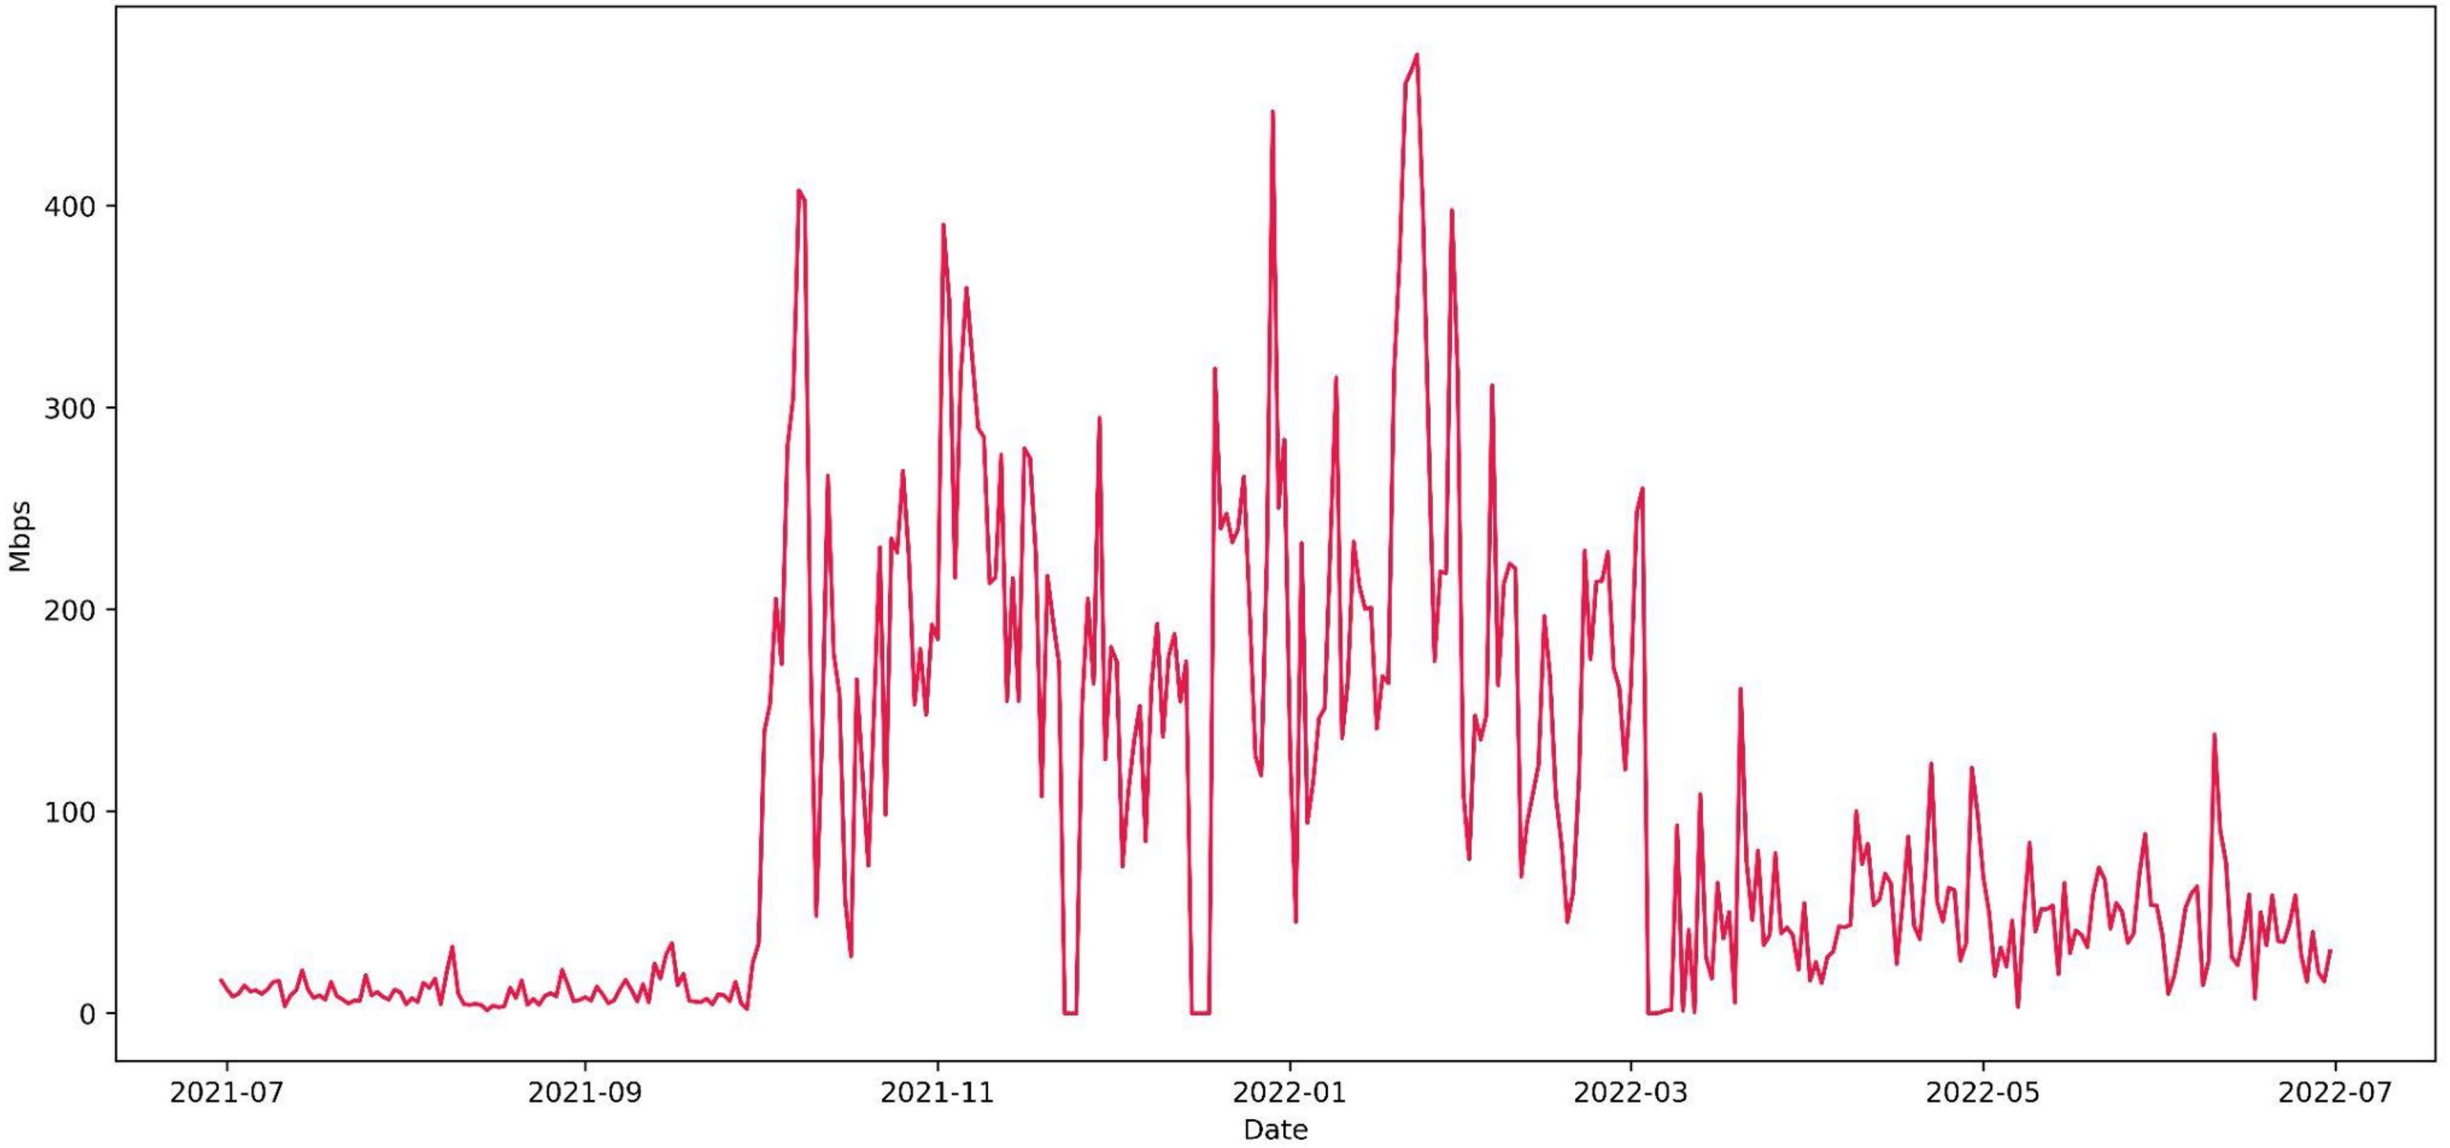

Daily cumulative throughput performance for cache miss

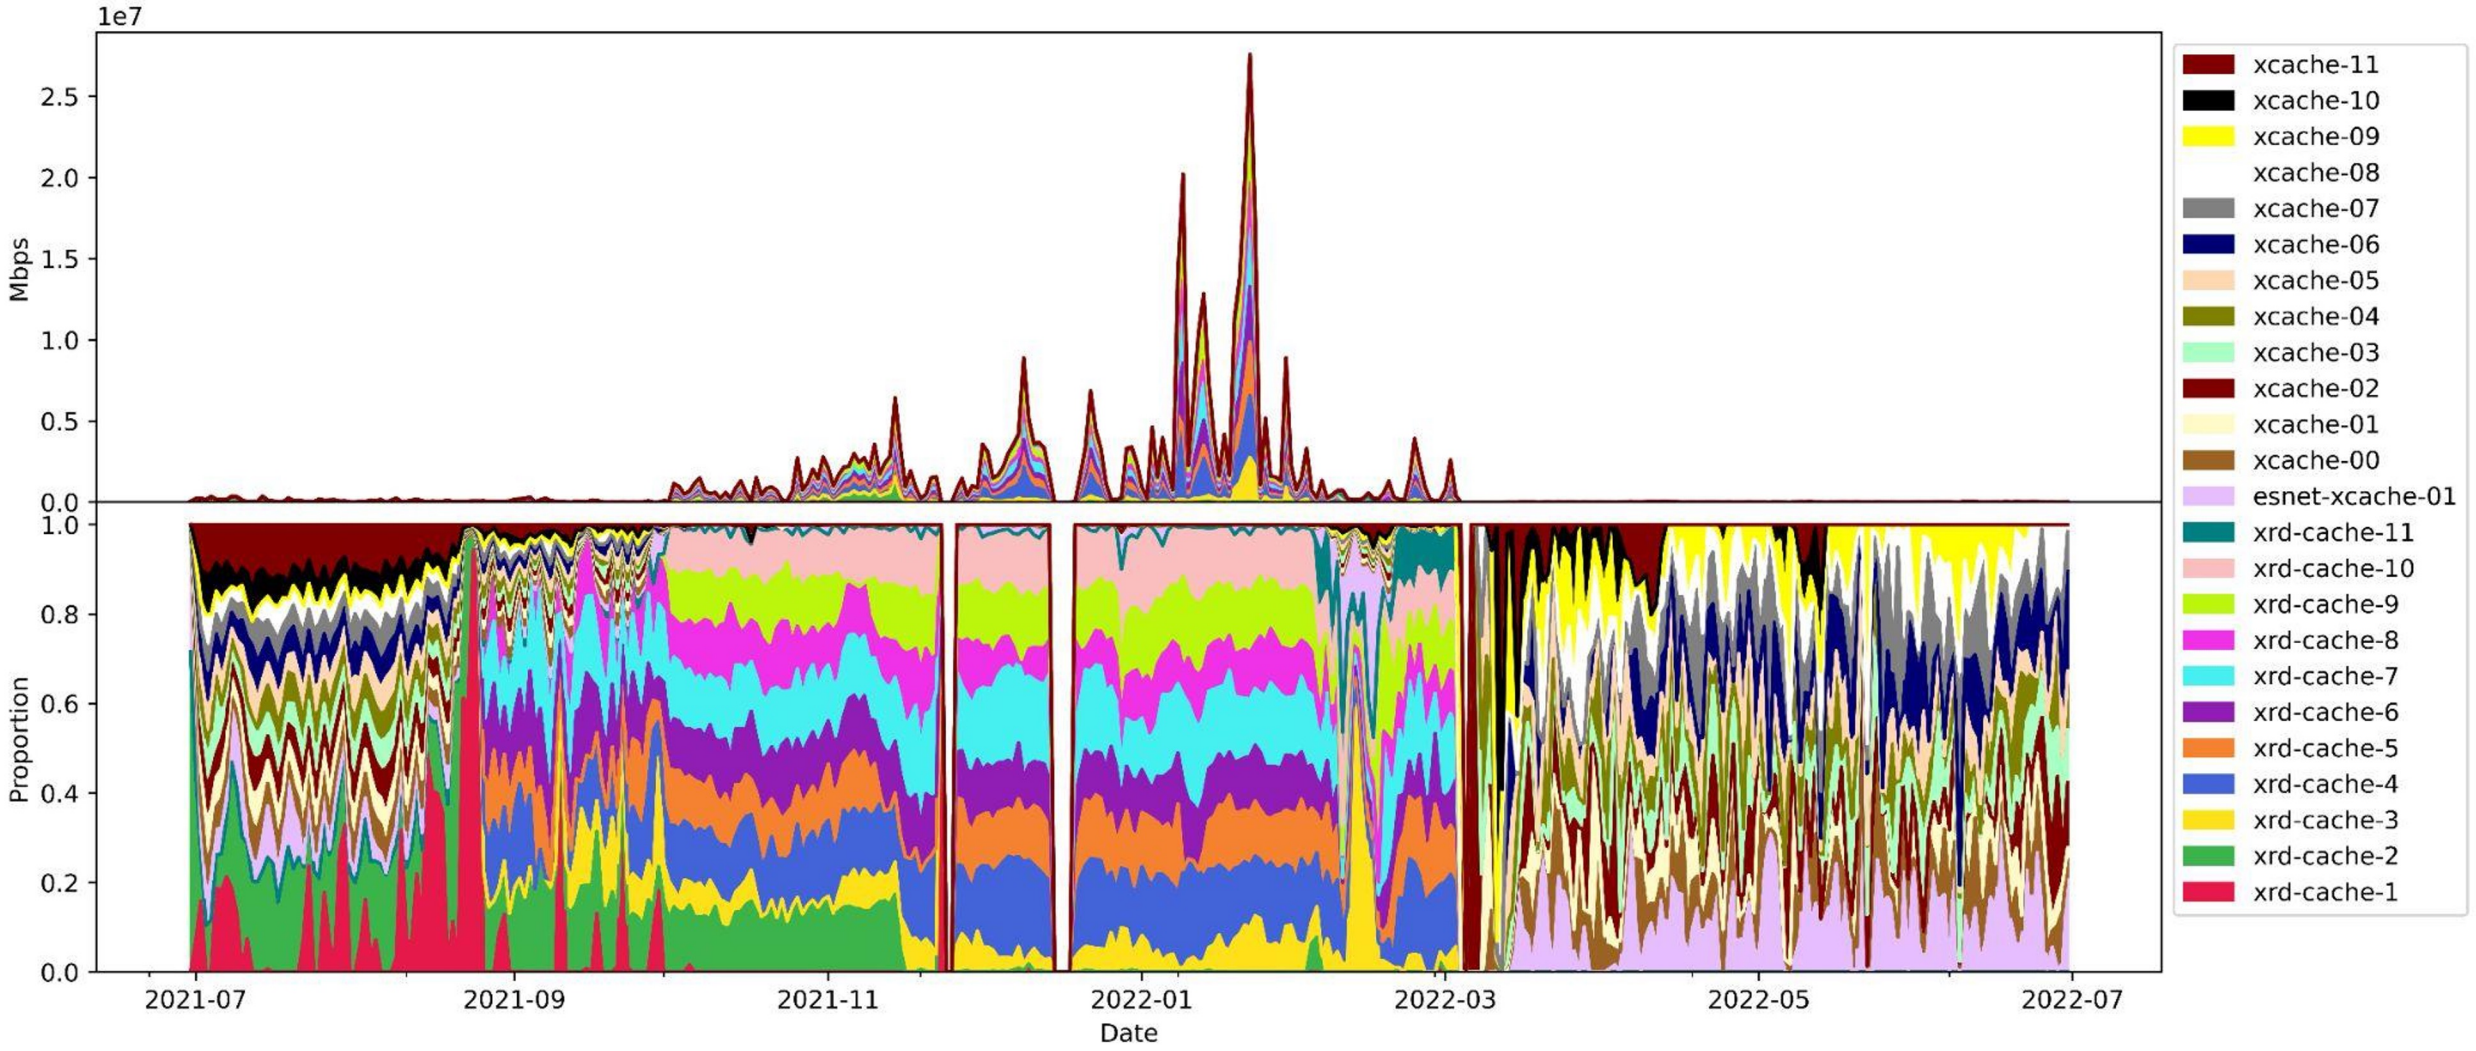

Hourly average data access throughput performance for cache hits per access

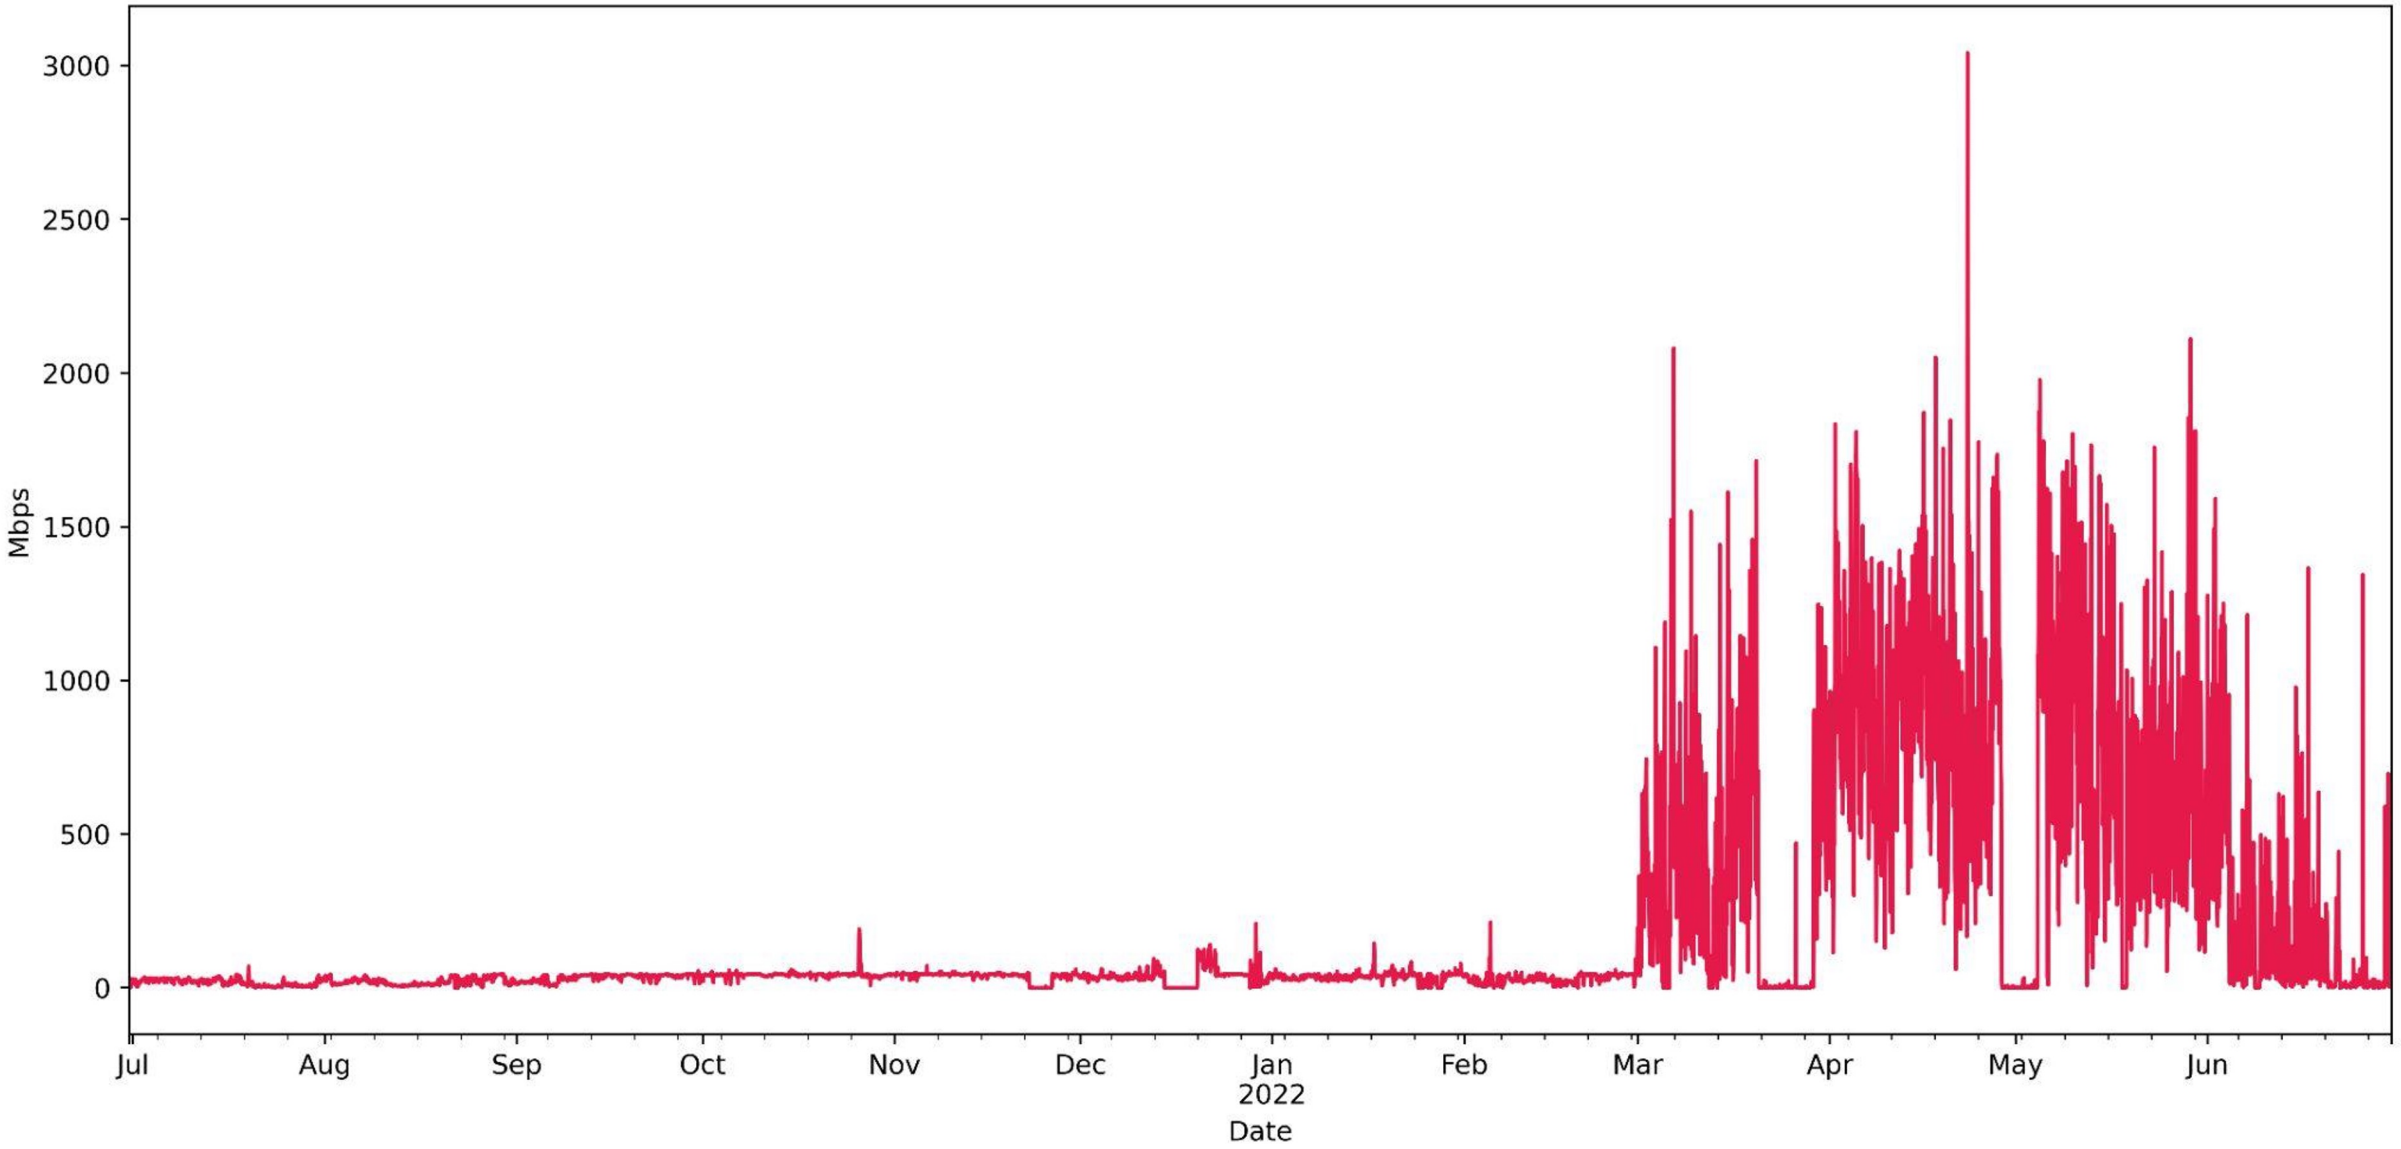

Hourly cumulative data access throughput performance for cache hits

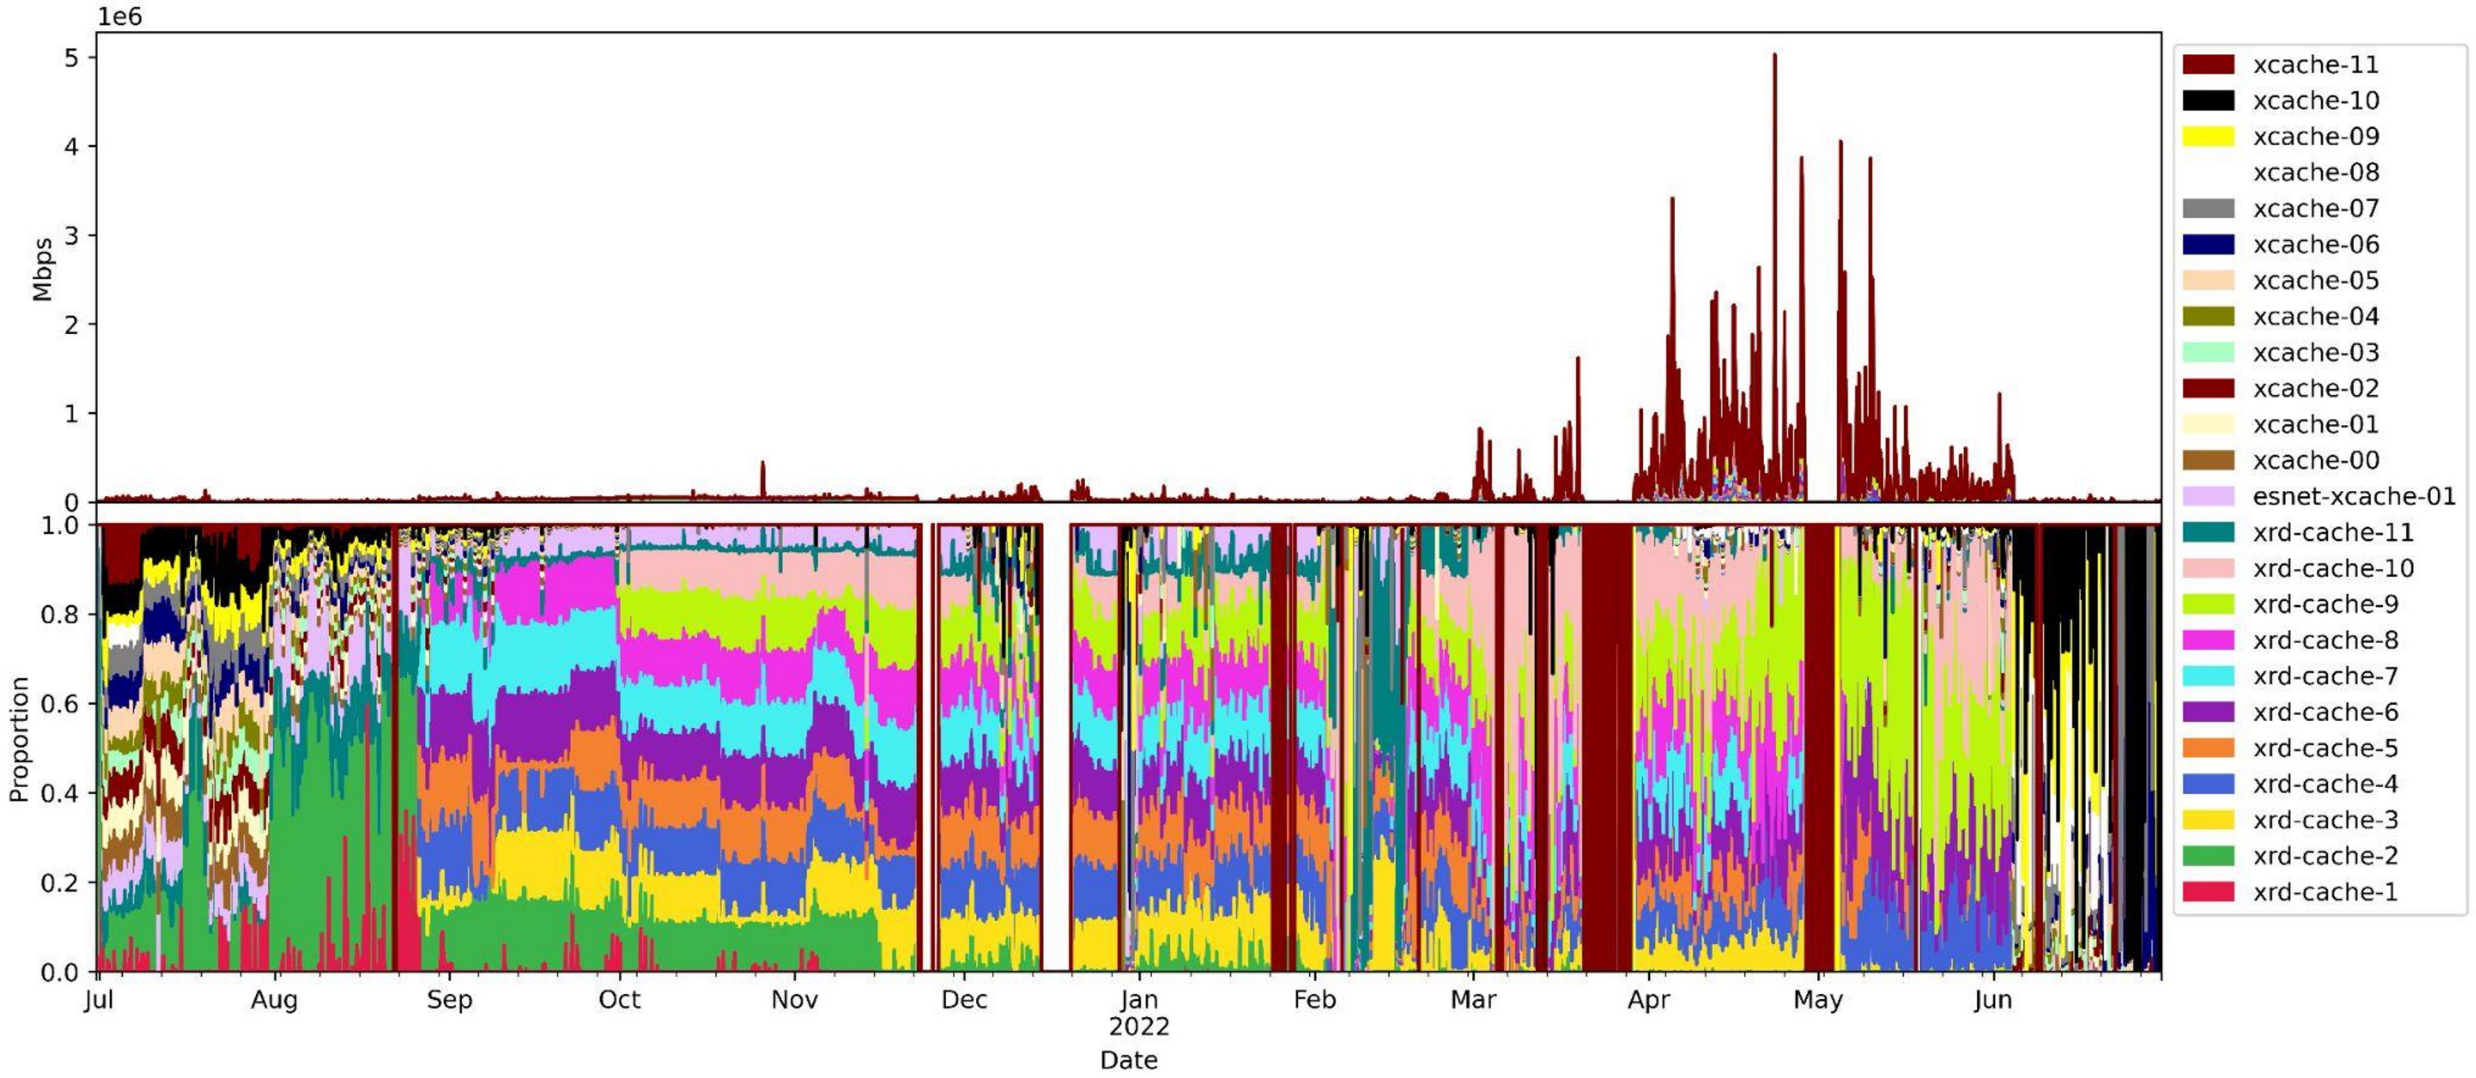

Hourly average data access throughput performance for cache miss per access

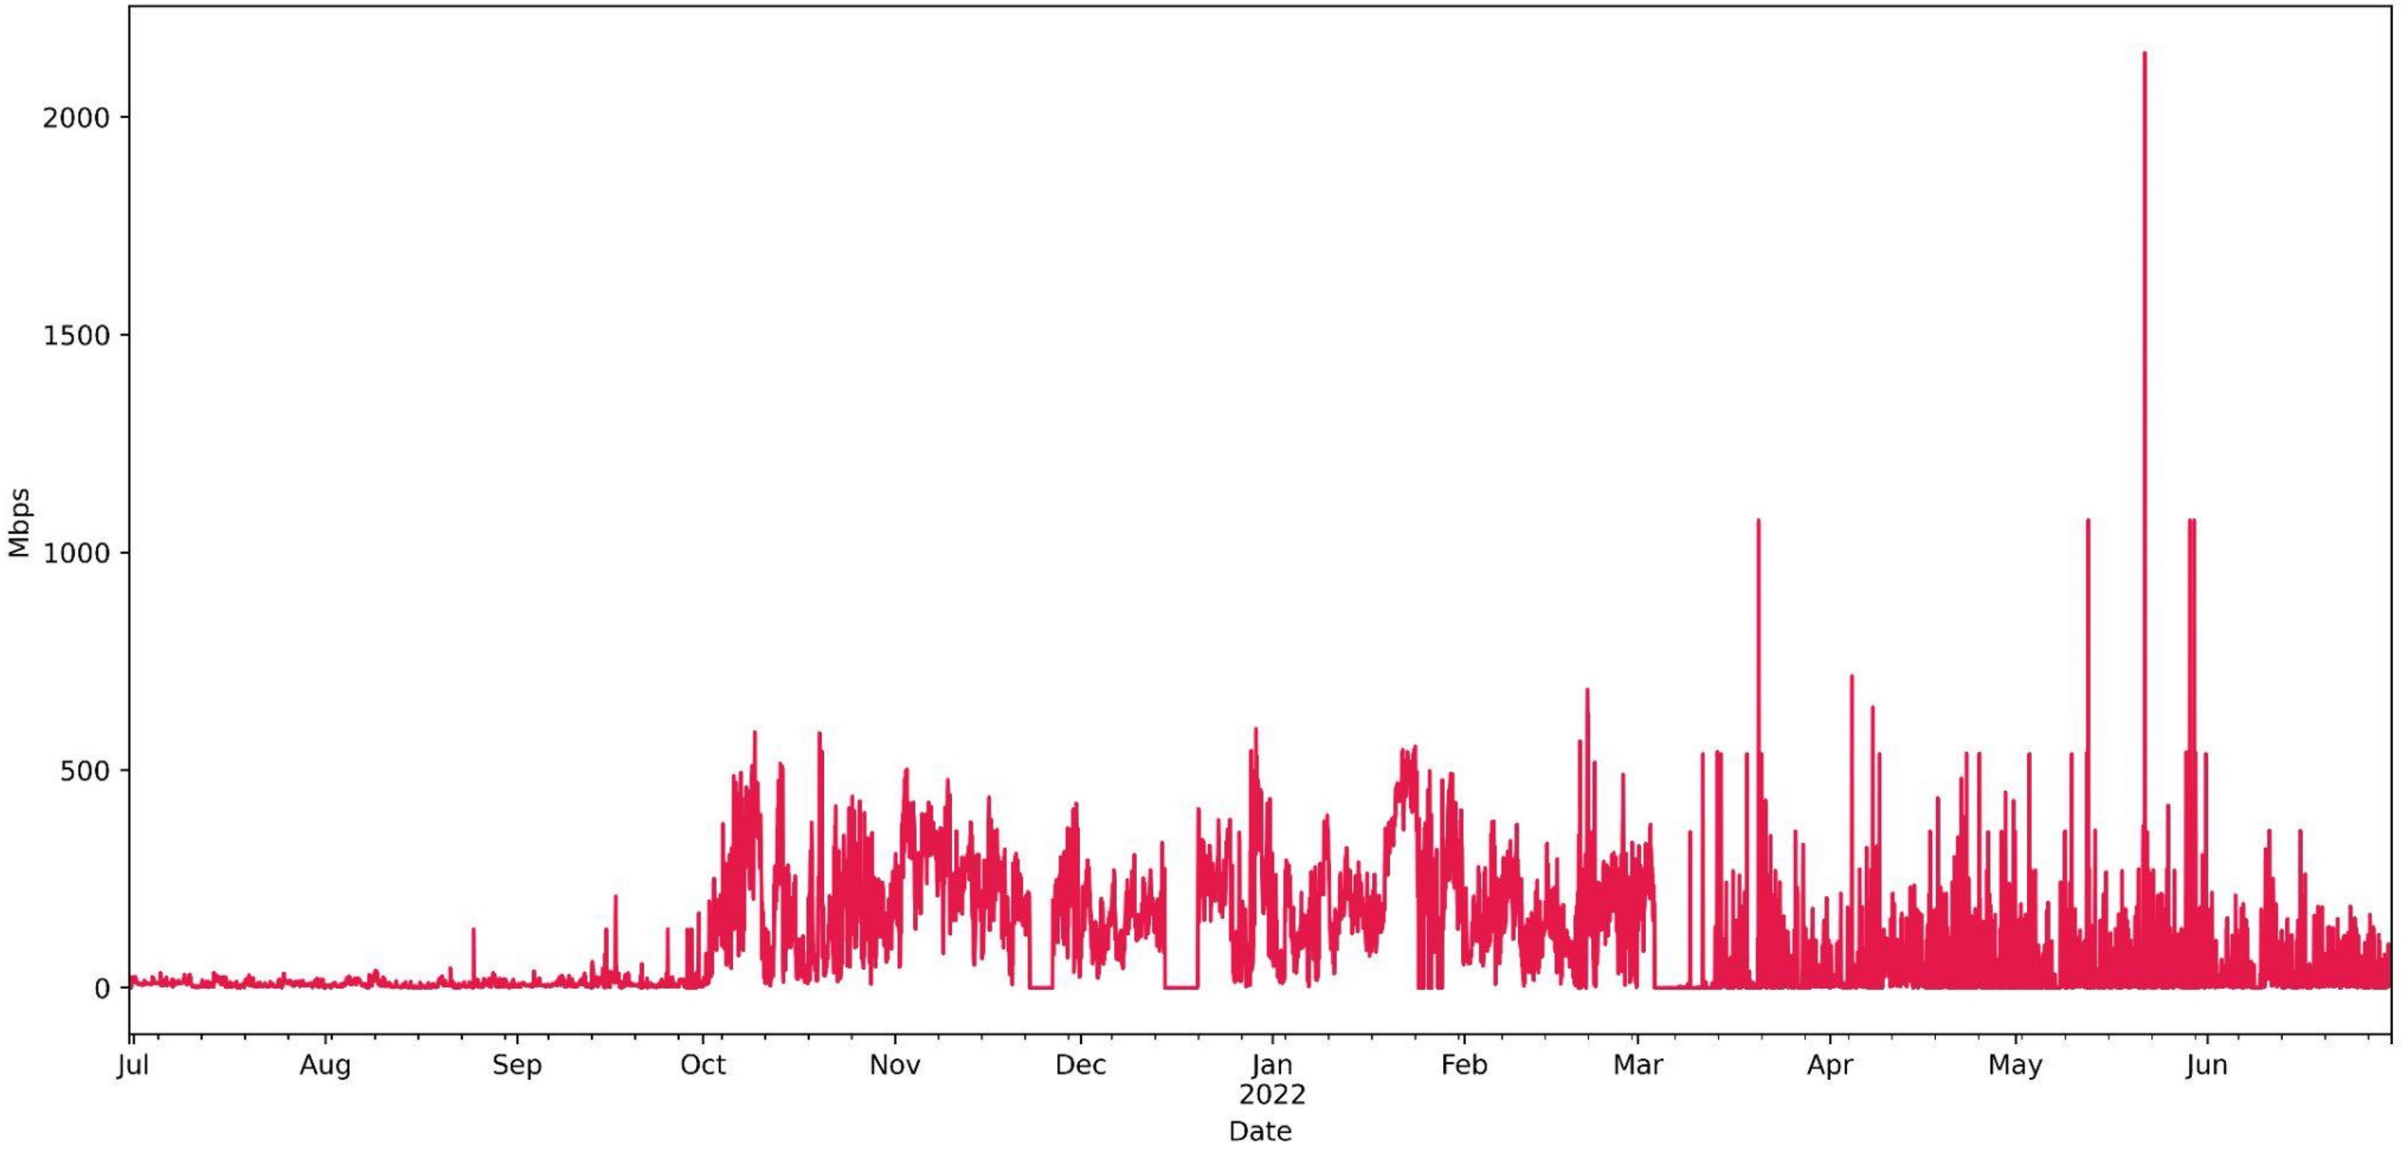

Hourly cumulative data access throughput performance for cache miss

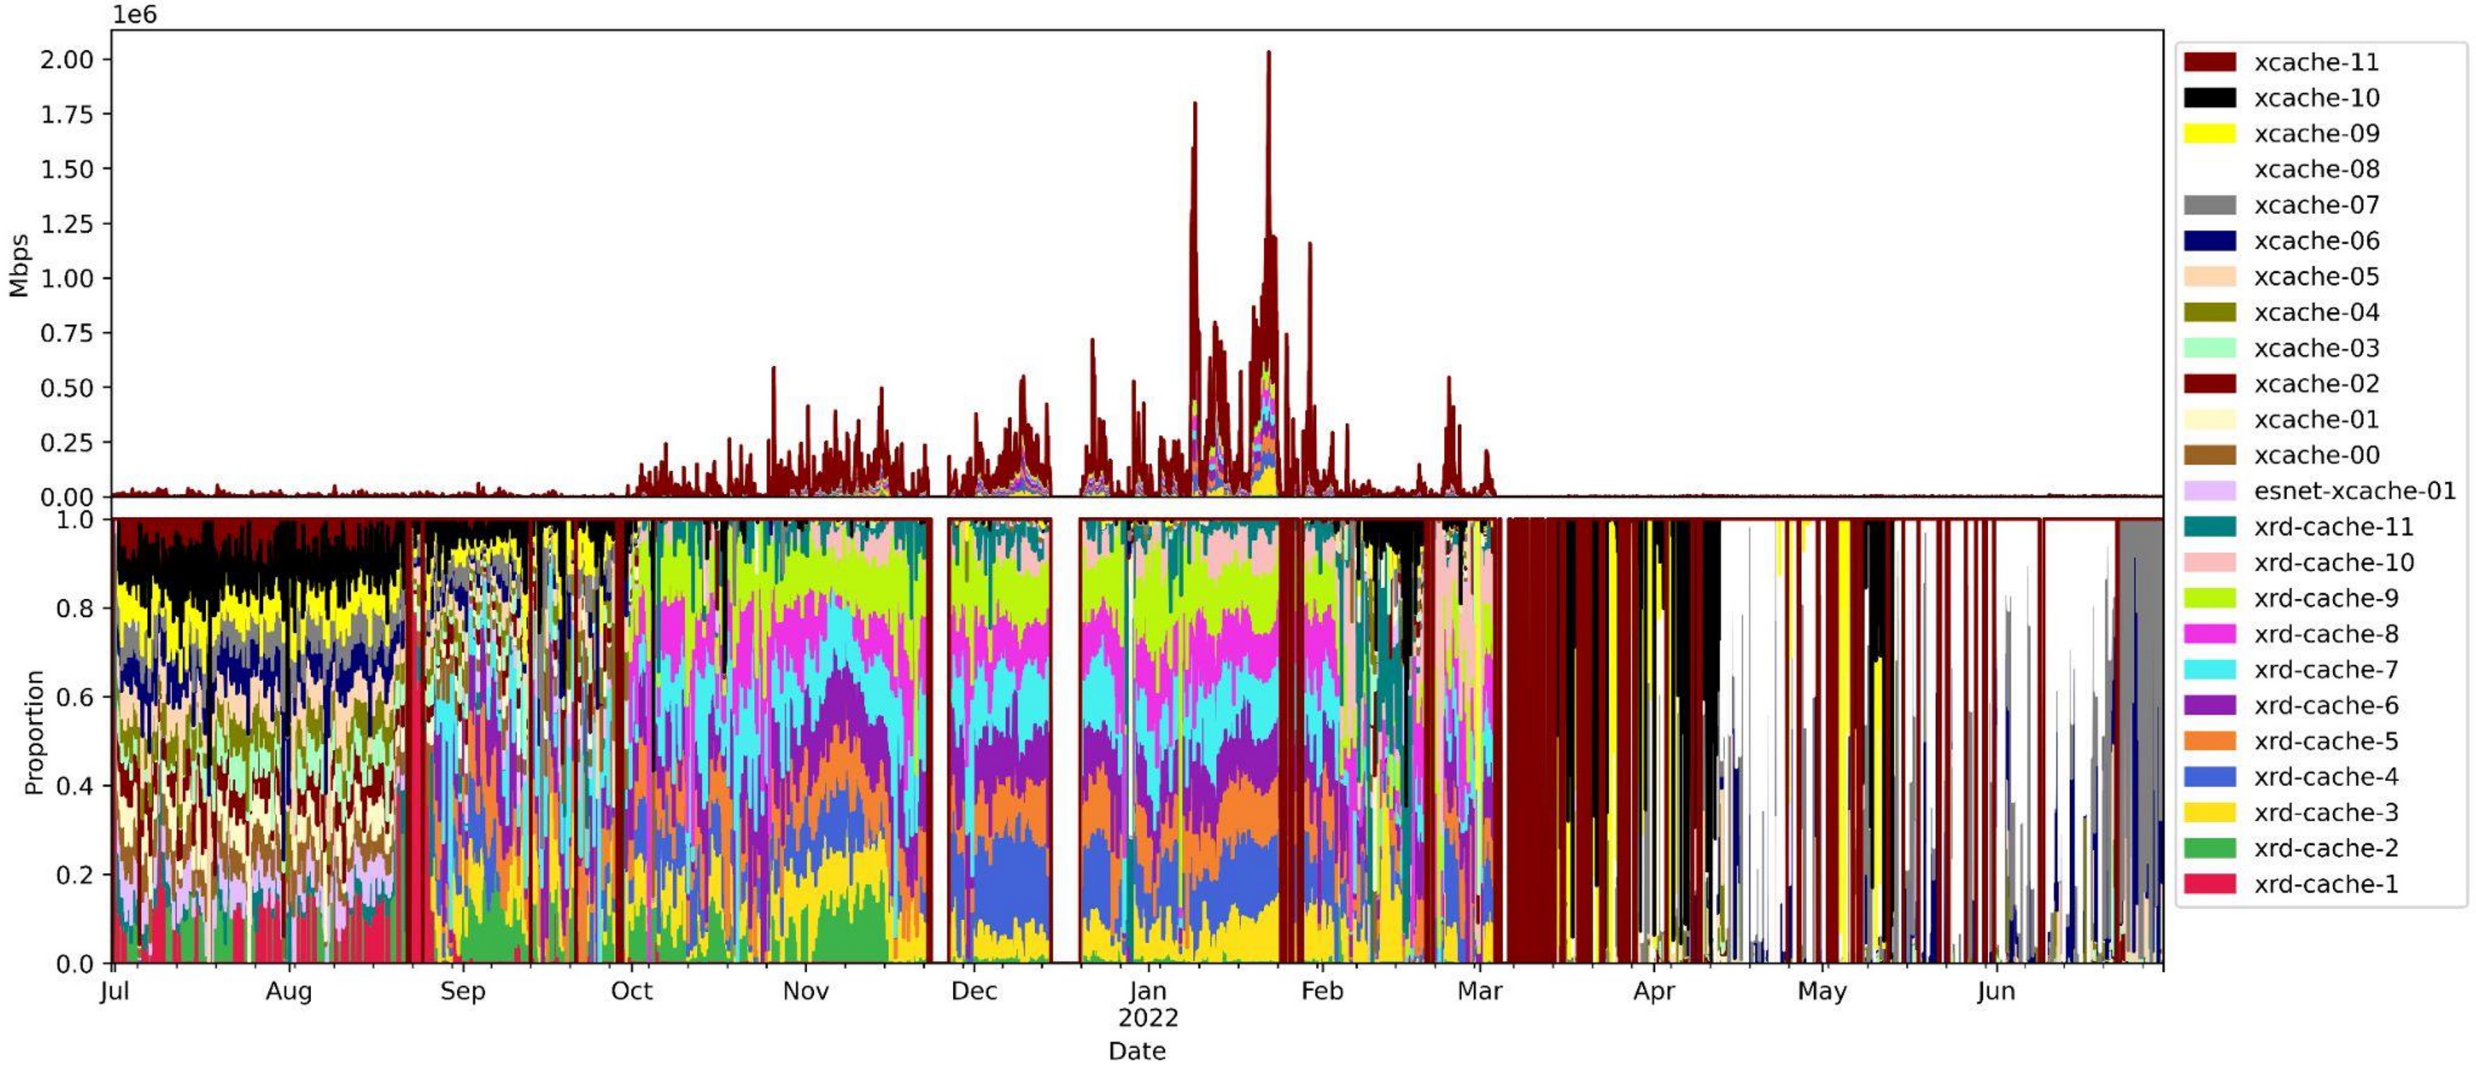

LSTM-based prediction on cache utilization

Model Construction

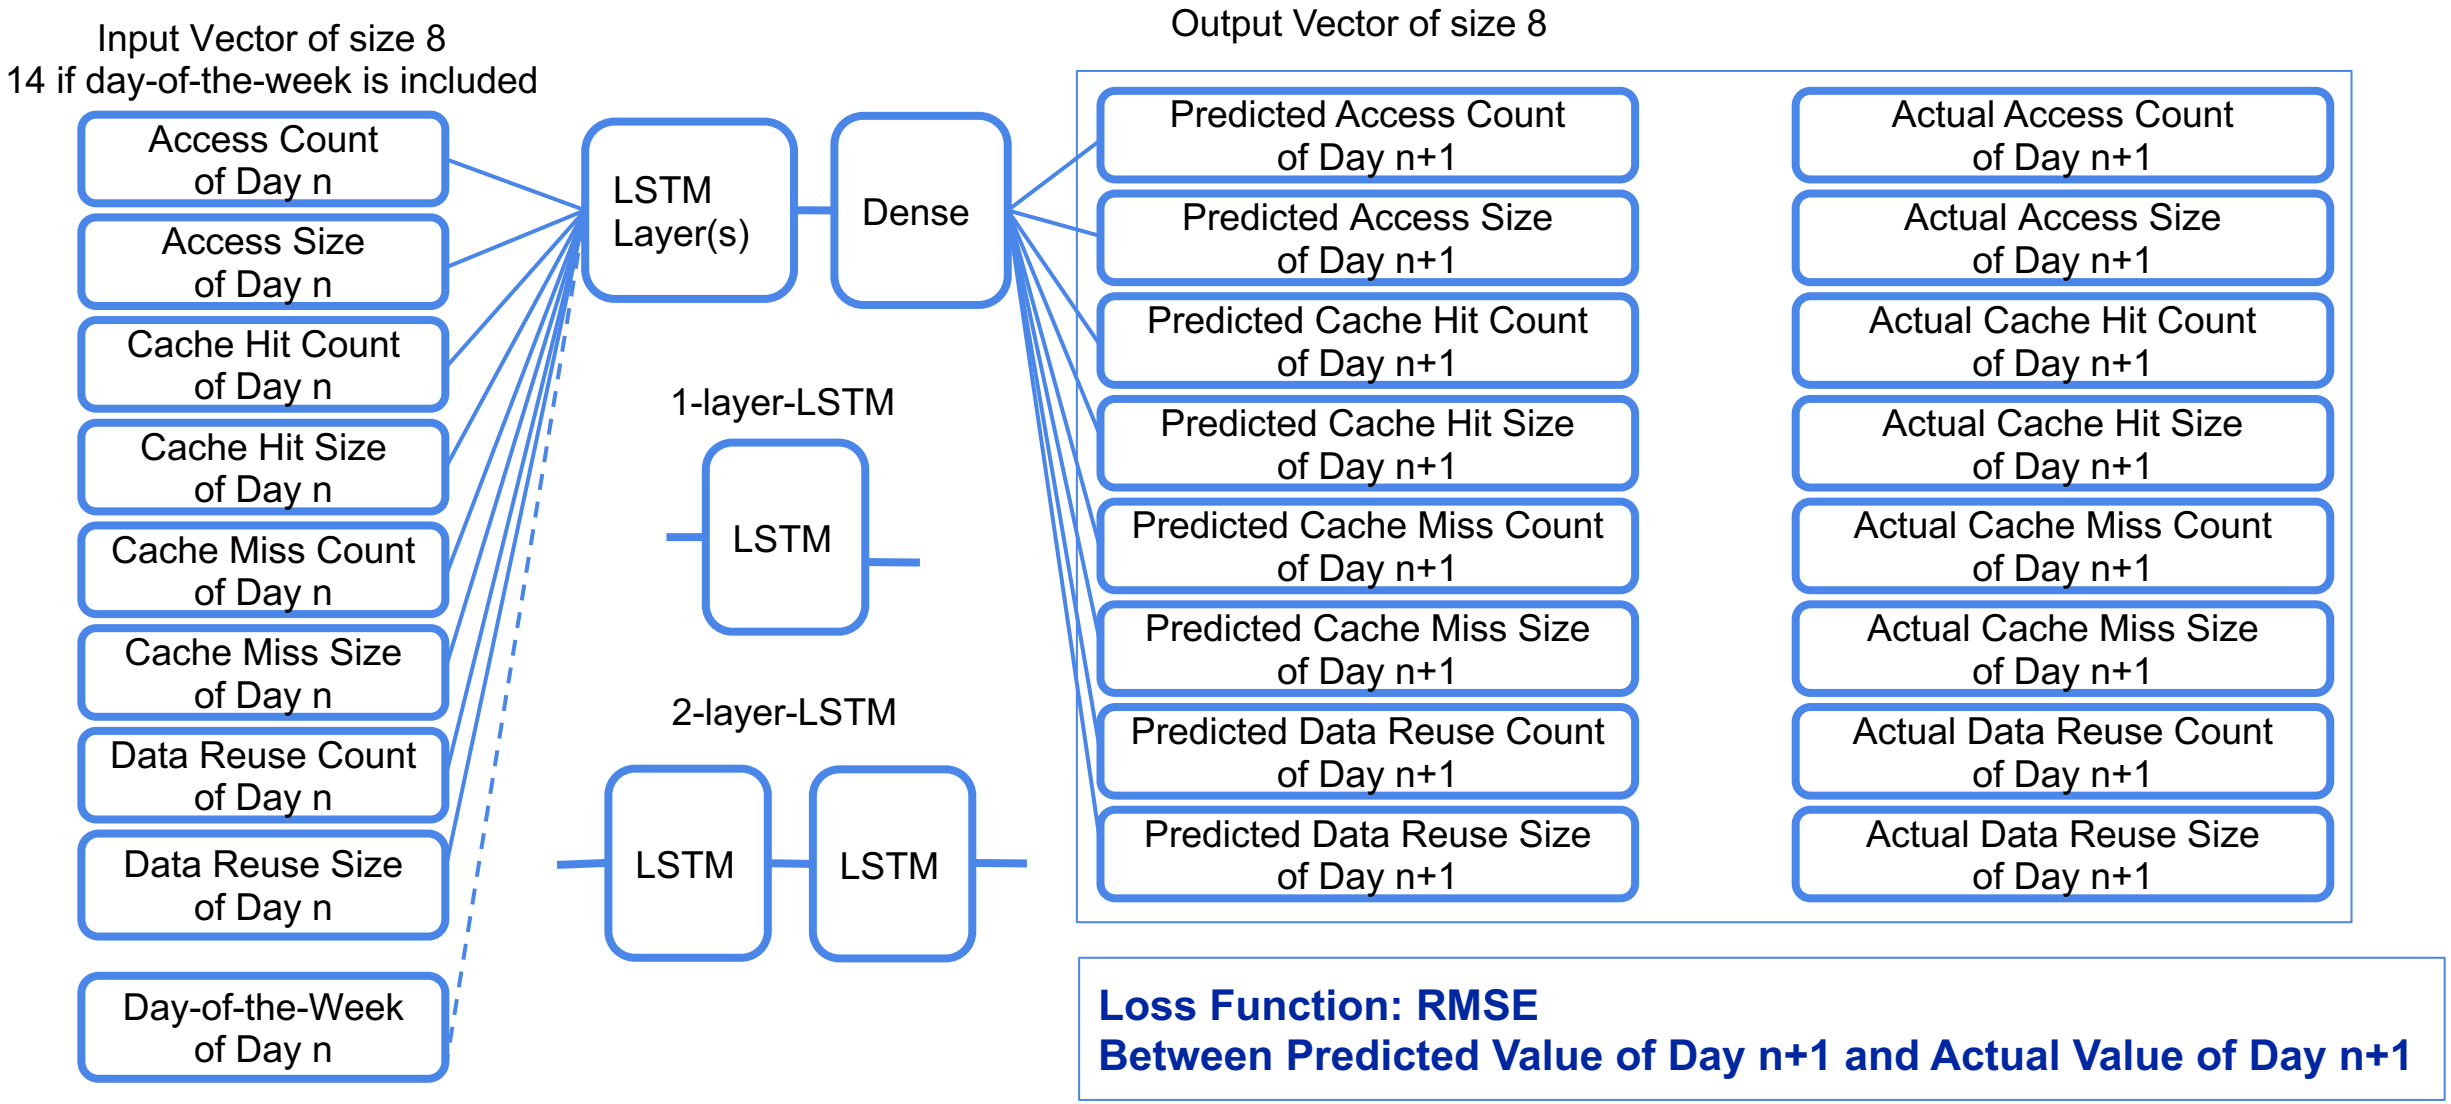

Loss Function: RMSE
Between Predicted Value of Day n+1 and Actual Value of Day n+1

Hyperparameters for daily LSTM model

Explored Parameters

Parameter	values
# of first layer LSTM unit	16, 32, 64, 128, 256
# of second layer LSTM unit	0, 16, 32, 64, 128, 256
first layer activation function	tanh, relu
second layer activation function	tanh, relu
dropout rate	0, 0.04, 0.1, 0.15
# of epchs	5, 10, 15, 25, 50, 75, 100

Model Parameter

	# of LSTM units	activation function	dropout rate	# of training epoch
values	128	tanh	0.04	50

Daily LSTM model results

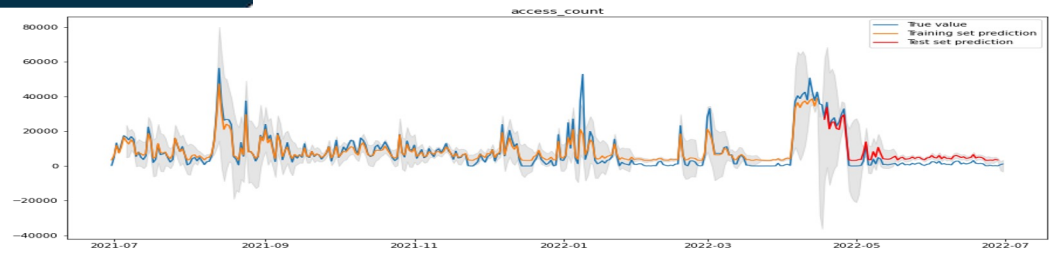

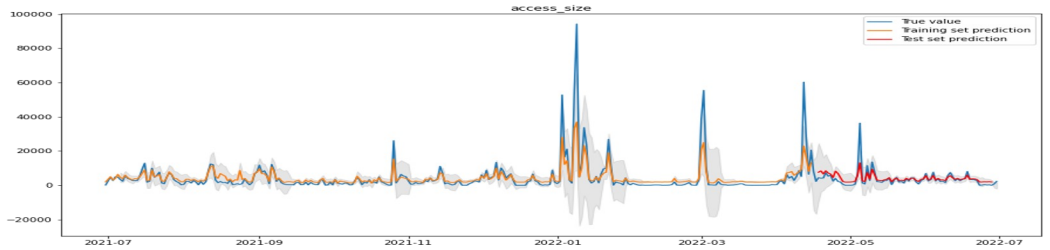

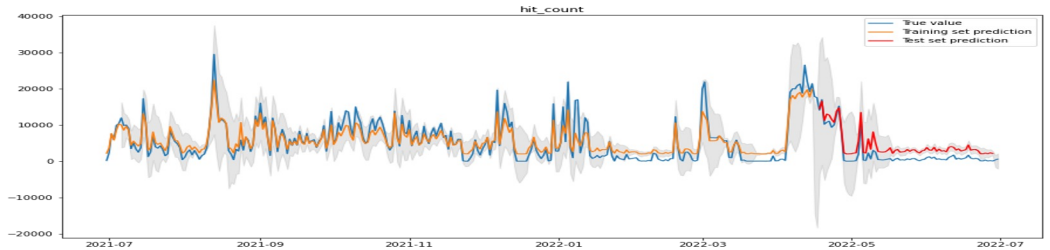

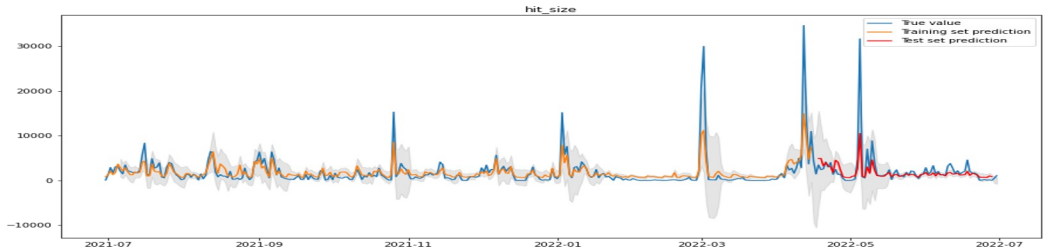

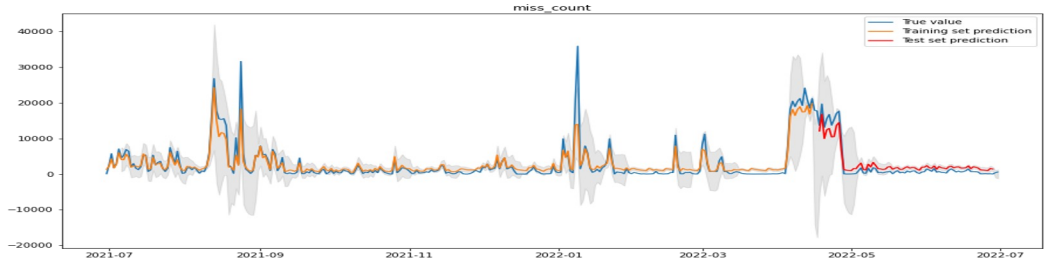

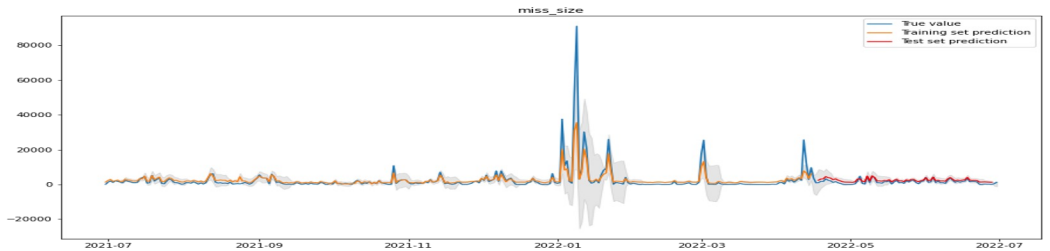

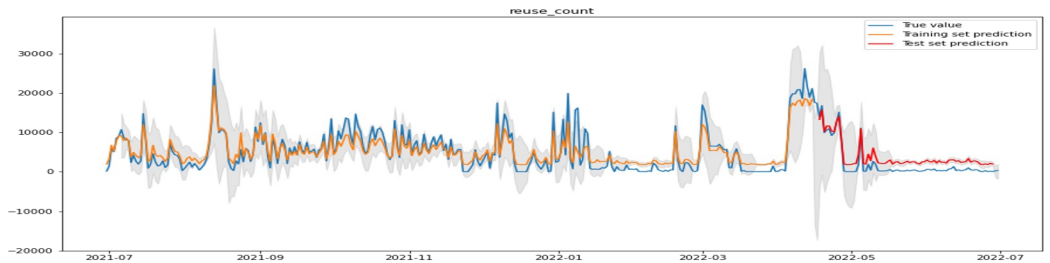

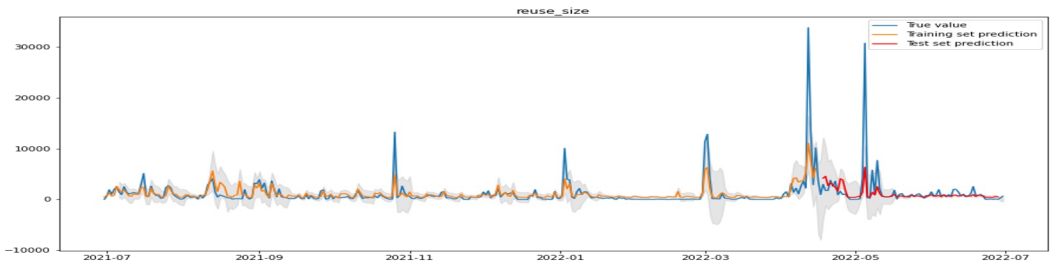

RMSE of daily LSTM model

	Without day-of-week		With day-of-week	
	Training RMSE	Test RMSE	Train RMSE	Test RMSE
Data Access Count	3797.97	4056.11	3436.02	3241.36
Data Access Volume	6048.17	3191.20	5234.59	3245.96
Cache Hit Count	2423.22	2335.64	2069.42	2225.02
Cache Hit Volume	2189.47	2958.92	2065.80	2676.81
Cache Miss Count	2687.23	2519.01	1984.61	1664.95
Cache Miss Volume	4873.69	1017.21	4020.90	1047.07
Data Re-use Count	2361.50	2222.56	2030.25	1960.07
Data Re-use Volume	1748.11	3091.21	1758.76	3104.17

Hyperparameters for daily LSTM model with 7 day moving average

- ## Explored Parameters

Parameter	values
# of first layer LSTM unit	16, 32, 64, 128, 256
# of second layer LSTM unit	0, 16, 32, 64, 128, 256
first layer activation function	tanh, relu
second layer activation function	tanh, relu
dropout rate	0, 0.04, 0.1, 0.15
# of epchs	5, 10, 15, 25, 50, 75, 100

- ## Model Parameters

	# of LSTM units	activation function	dropout rate	# of training epoch
values	64	relu	0.1	50

Daily LSTM model results with 7 day moving average

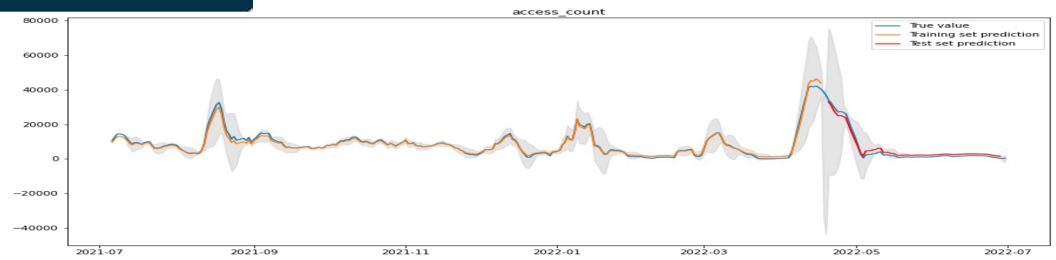

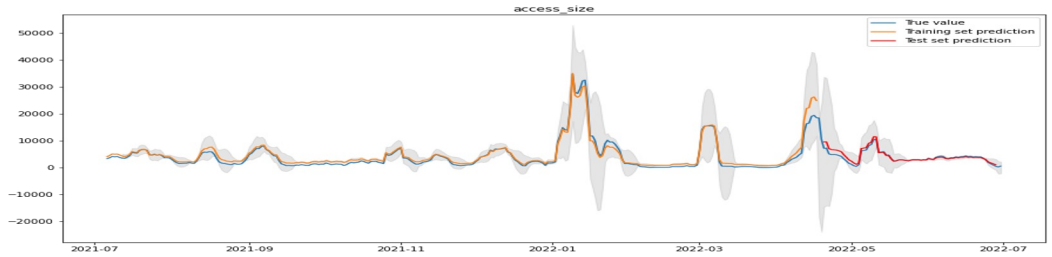

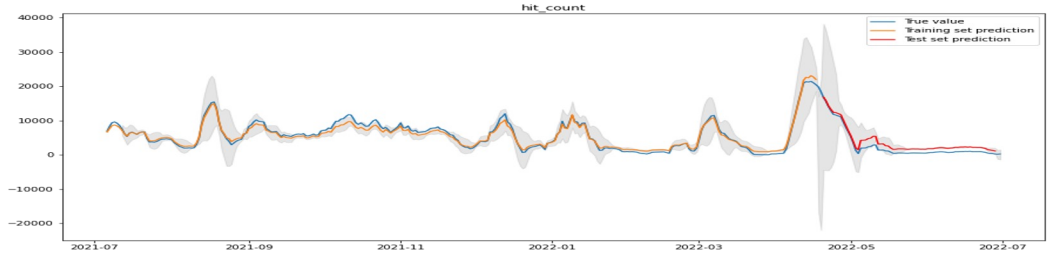

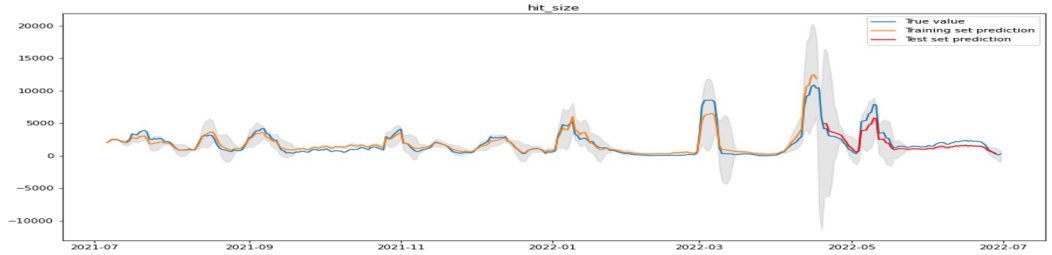

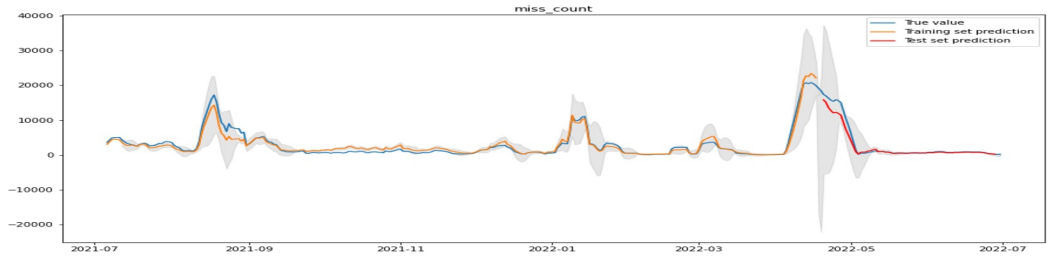

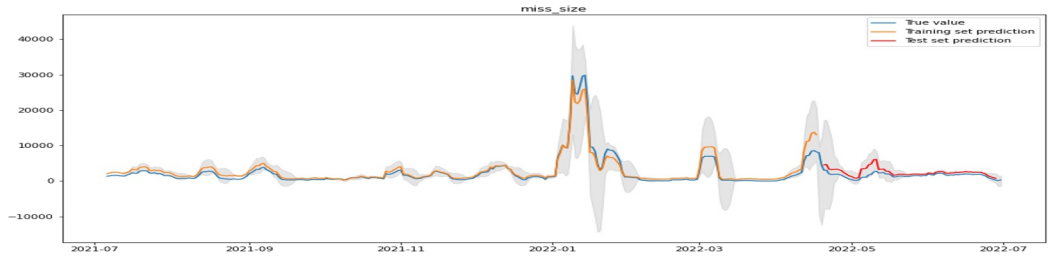

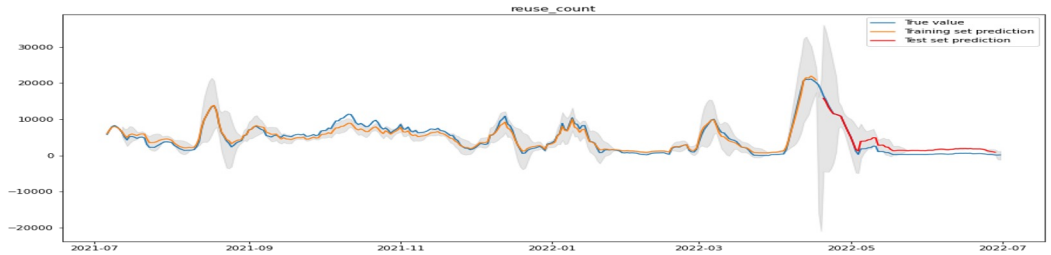

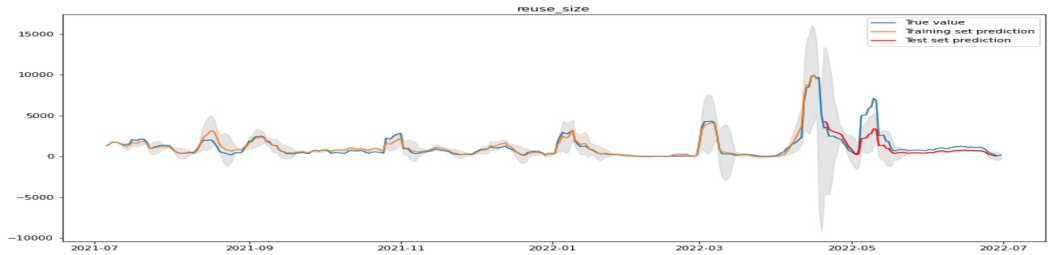

Daily LSTM model results with 7 day moving average

	Train RMSE of MA LSTM model	Test RMSE of MA LSTM model
Access Count	1002.02	1452.05
Access Size	1221.59	831.26
Cache Hit Count	727.46	1318.12
Cache Hit Size	542.65	773.34
Cache Miss Count	990.86	1255.52
Cache Miss Size	1102.32	1165.82
Data Reuse Count	820.27	1259.19
Data Reuse Size	294.74	1119.77

Hyperparameters for hourly LSTM model

Explored Parameters

Parameter	values
# of first layer LSTM unit	16, 32, 64, 128, 256
# of second layer LSTM unit	0, 16, 32, 64, 128, 256
first layer activation function	tanh, relu
second layer activation function	tanh, relu
dropout rate	0, 0.04, 0.1, 0.15
# of epchs	5, 10, 15, 25, 50, 75, 100

Model Parameter

	# of LSTM units	activation function	dropout rate	# of training epoch
values	128	relu	0.1	50

Hourly LSTM model results

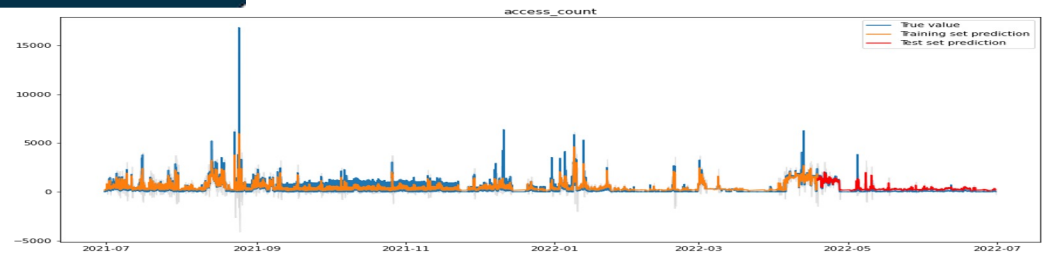

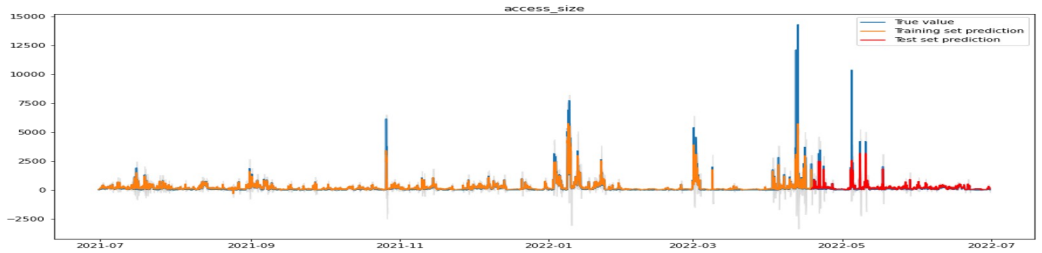

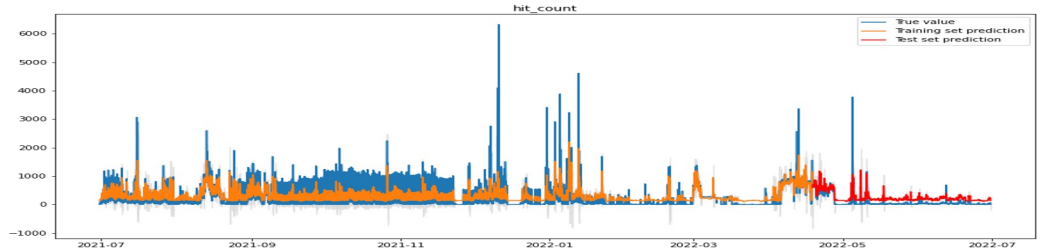

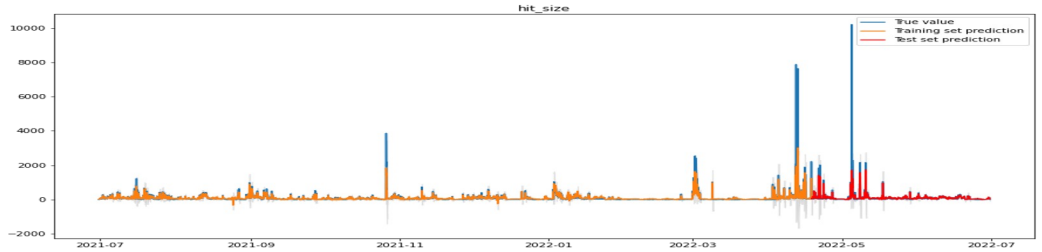

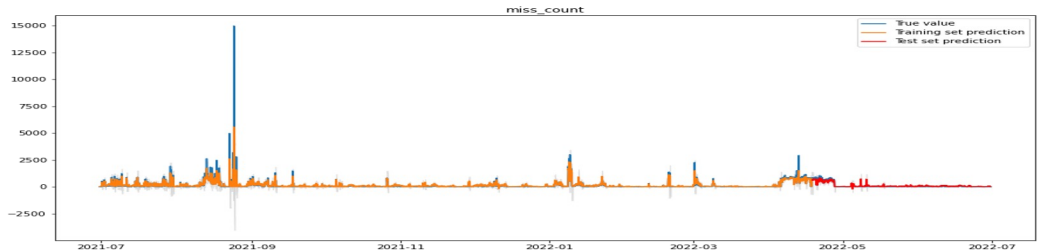

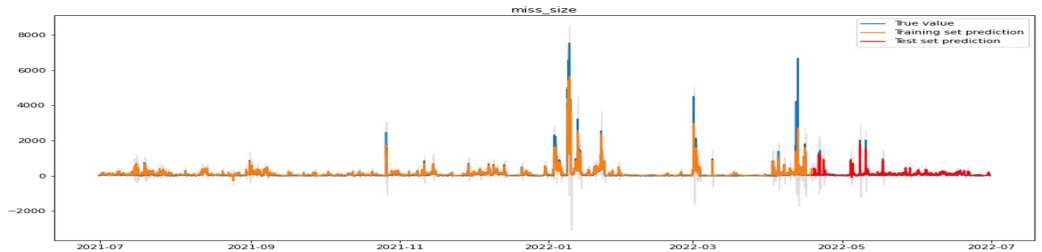

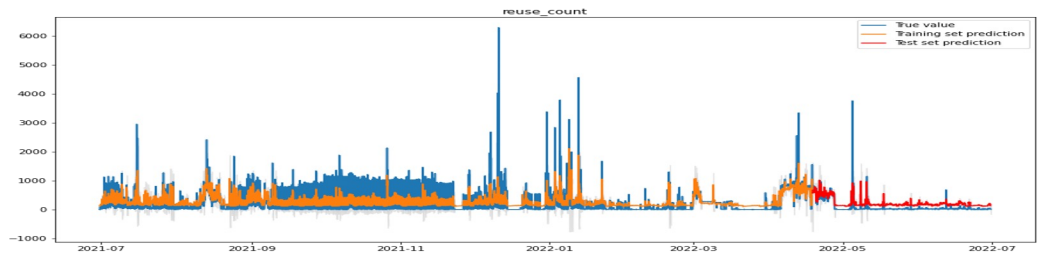

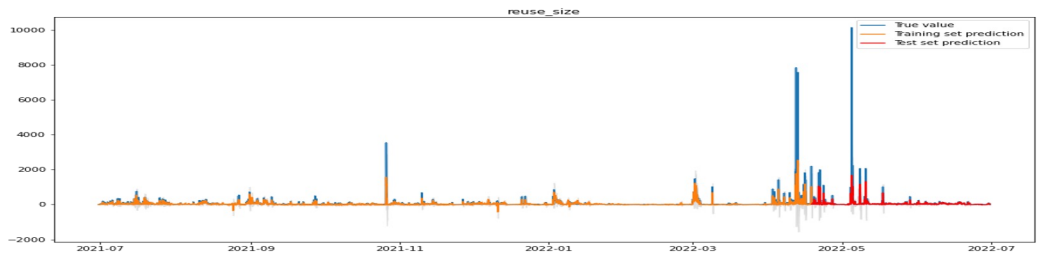

RMSE of hourly LSTM model

	Without day-of-week		With day-of-week	
	Training RMSE	Test RMSE	Train RMSE	Test RMSE
Data Access Count	253.92	165.41	301.76	216.21
Data Access Volume	192.70	212.22	218.72	239.22
Cache Hit Count	234.28	137.63	252.78	176.83
Cache Hit Volume	118.38	200.27	127.45	235.78
Cache Miss Count	141.11	43.89	141.53	52.61
Cache Miss Volume	90.99	54.07	105.19	50.38
Data Re-use Count	236.15	128.55	252.93	163.05
Data Re-use Volume	120.72	208.00	131.37	235.05

LSTM-based prediction on data access throughput performance

Model Construction

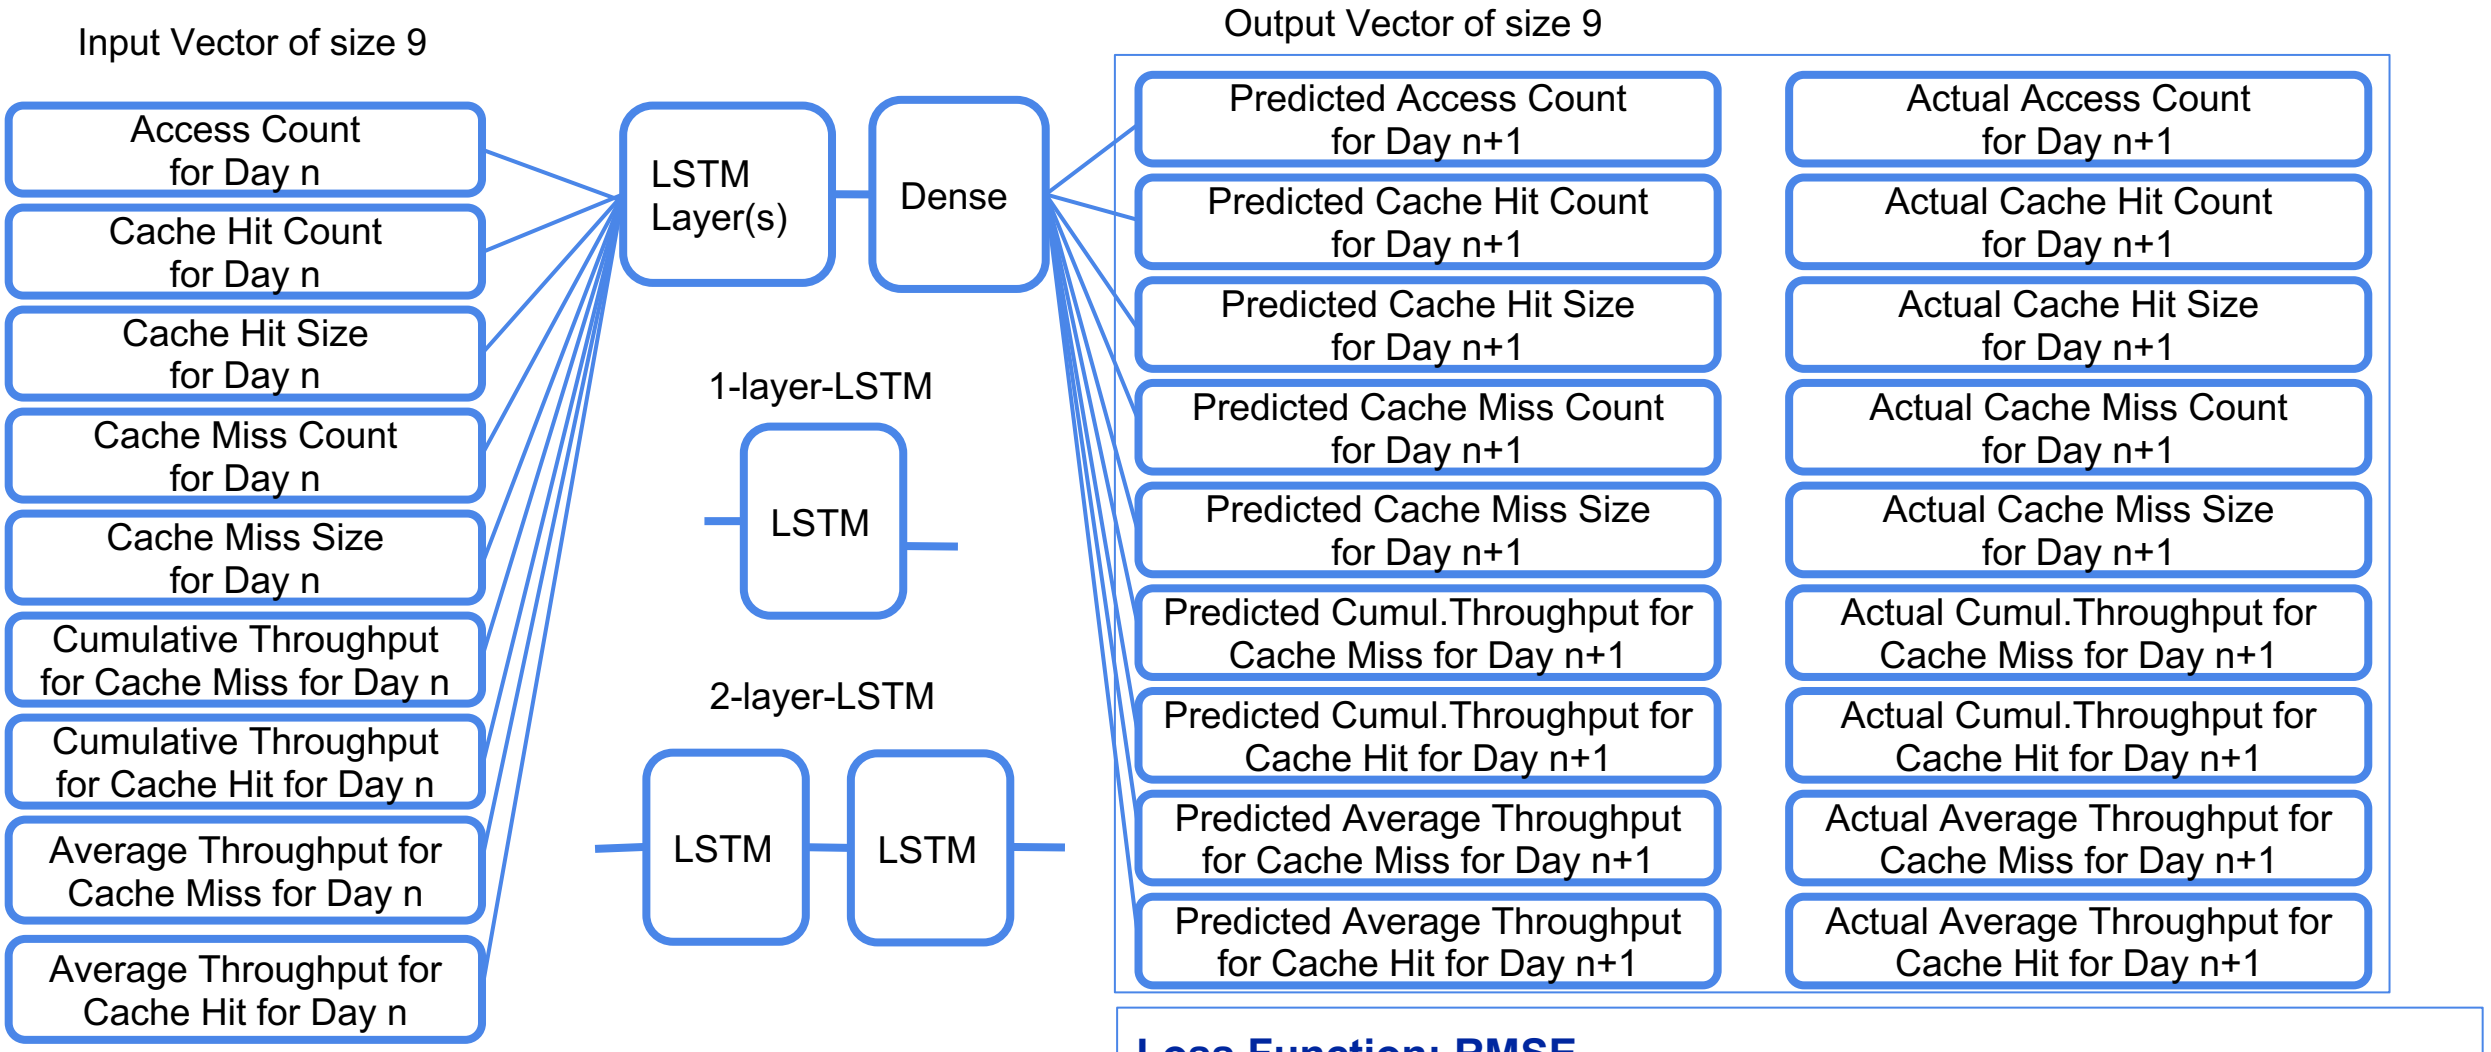

Loss Function: RMSE
Between Predicted Value of Day n+1 and Actual Value of Day n+1

Hyperparameters for LSTM model for data access throughput performance

	# of LSTM units	activation function	dropout rate	# of training epoch
Daily throughput	128	tanh	0.04	50
Daily throughput with moving average	64	relu	0.1	50
Hourly throughput	128	tanh	0.04	50
Hourly throughput with moving average	64	relu	0.1	50

Daily LSTM model results for data access throughput performance

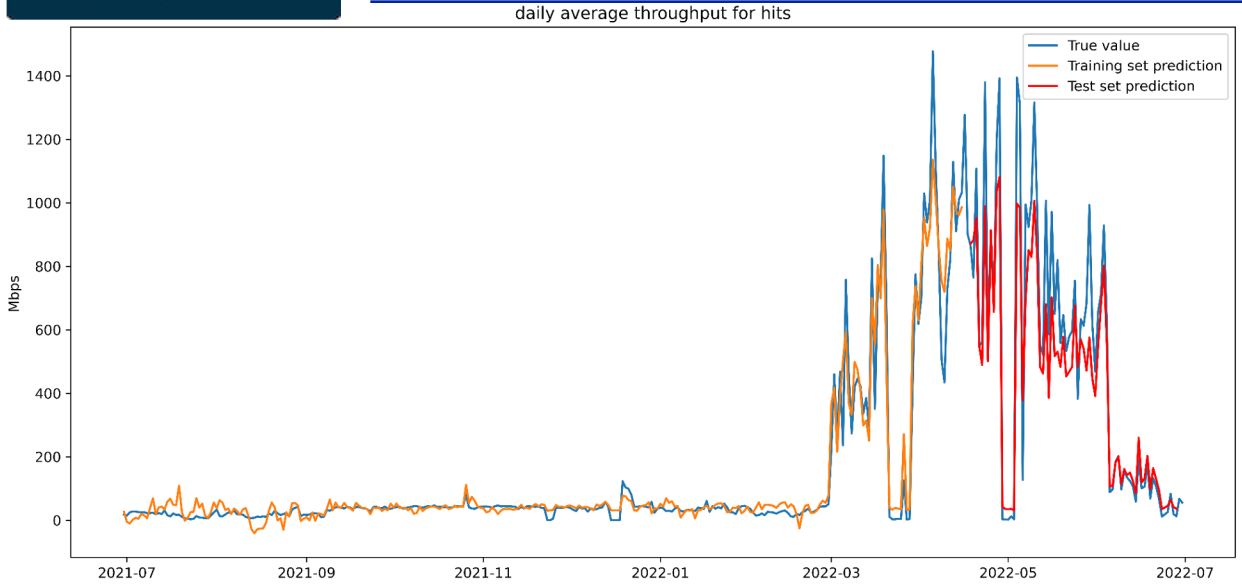

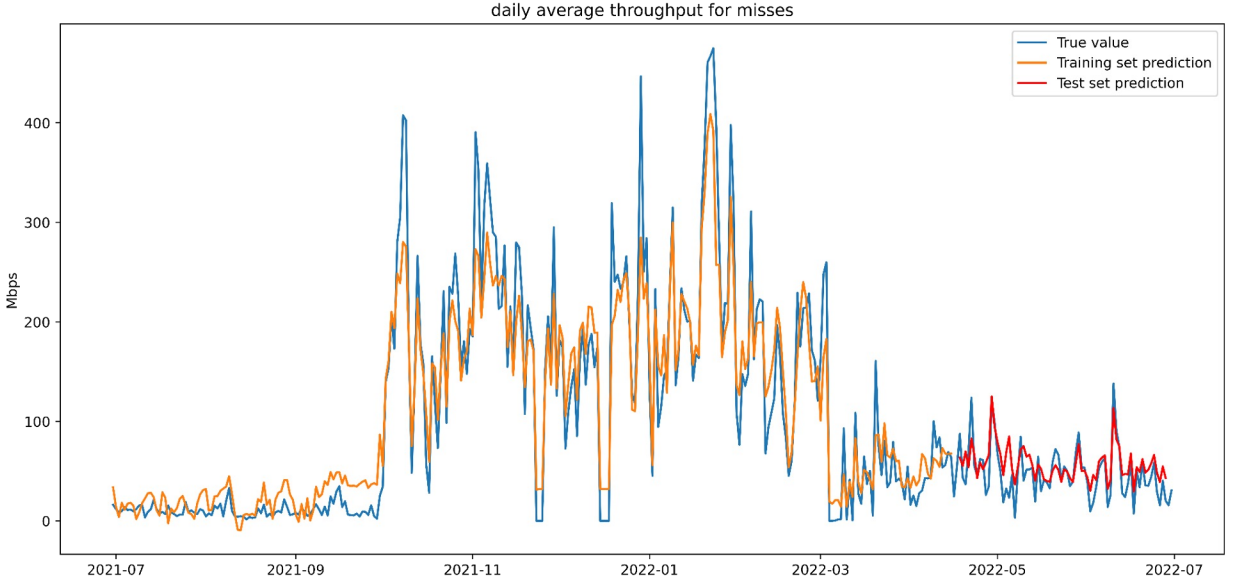

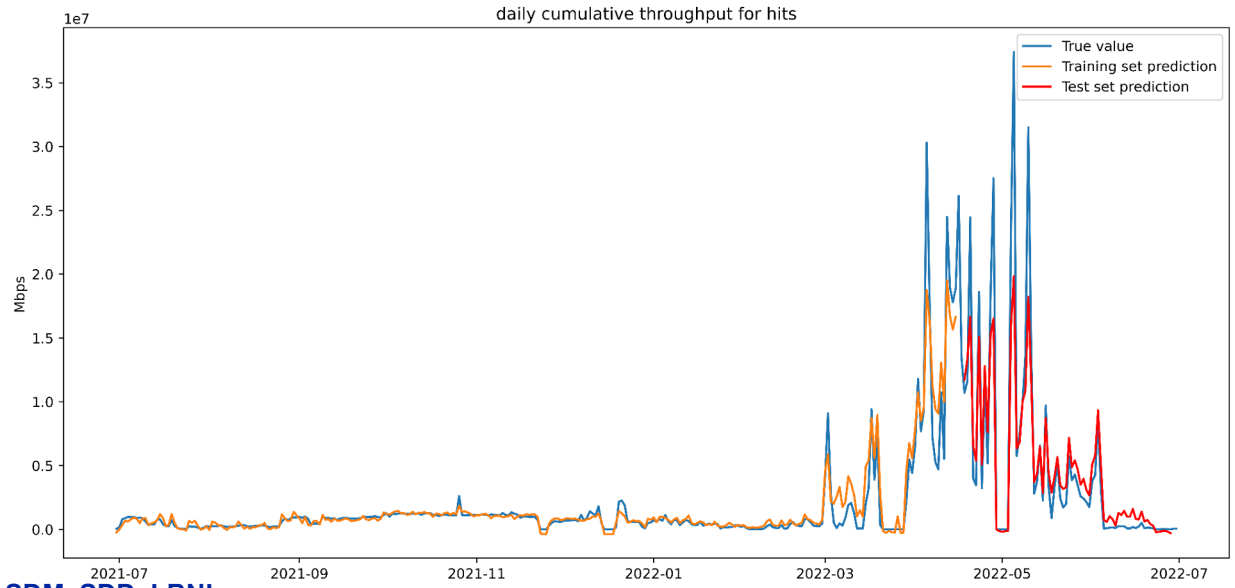

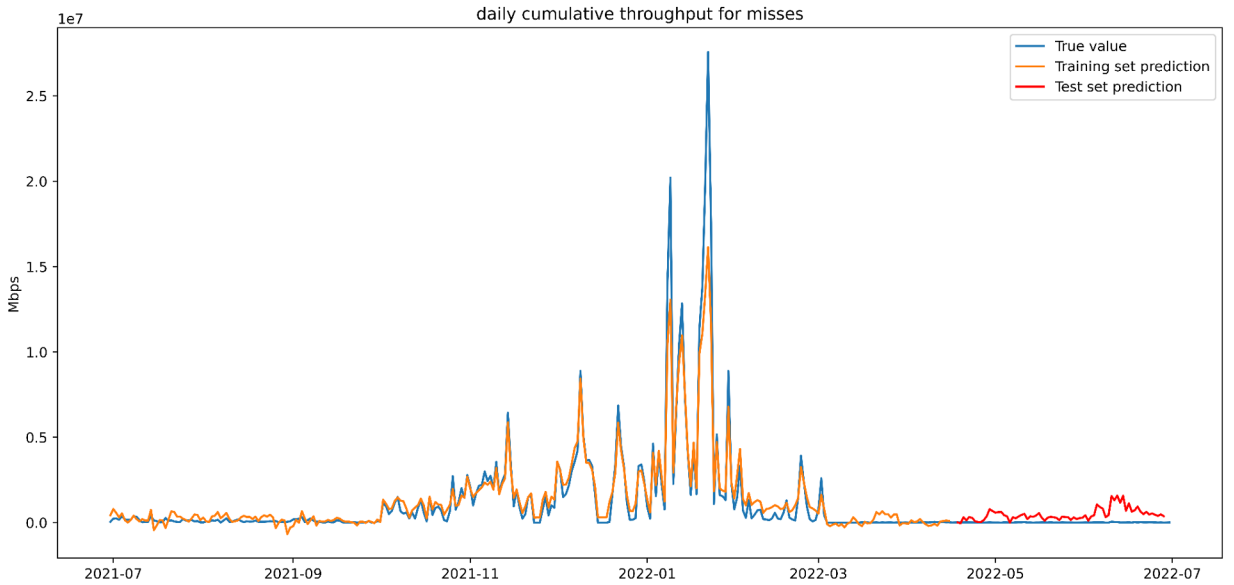

RMSE of Daily LSTM model results for data access throughput performance

	Train RMSE	Test RMSE
Daily average throughput for cache misses	34.57	18.27
Daily average throughput for cache hits	52.29	149.06
Daily cumulative throughput for cache misses	1068766.00	582794.20
Daily cumulative throughput for cache hits	1096641.00	3392070.00

Daily LSTM model results with 7-day moving average for data access throughput performance

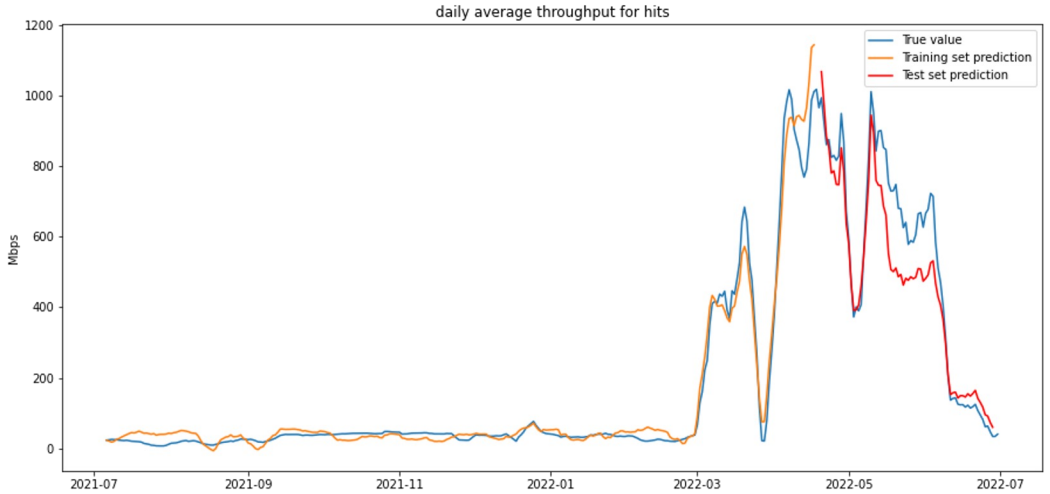

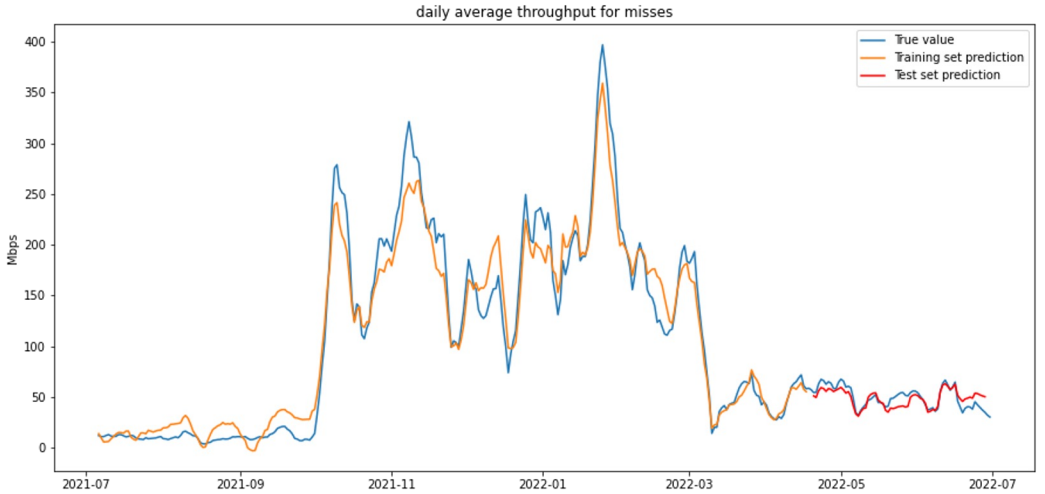

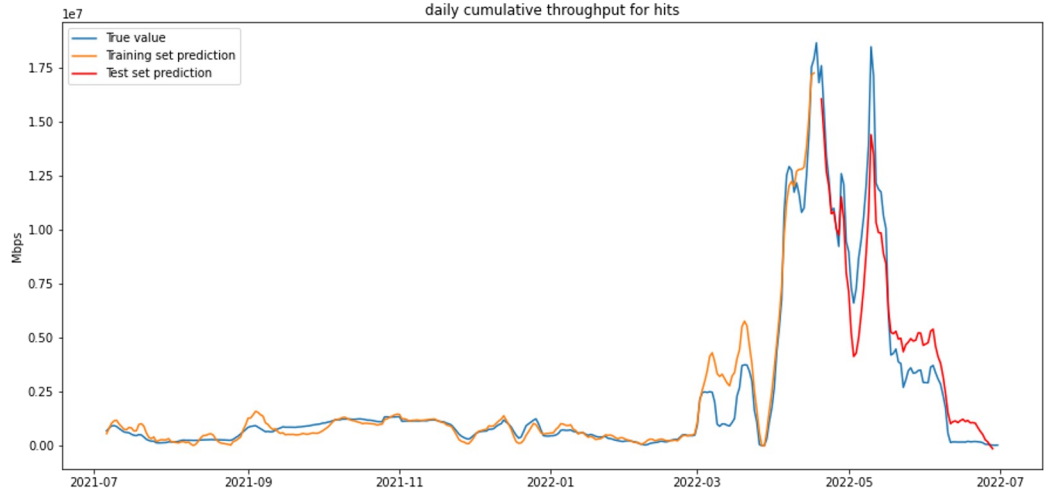

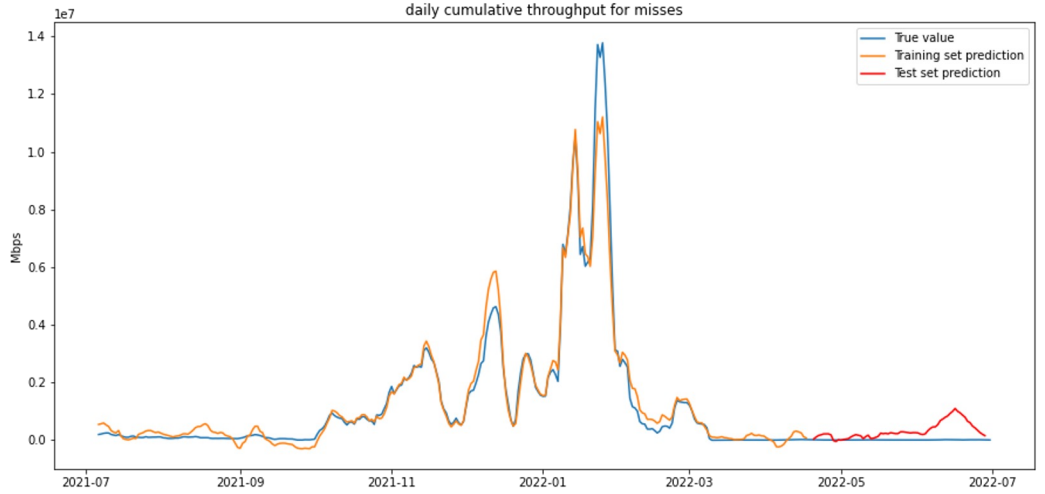

RMSE of Daily LSTM model results with 7 day MA for data access throughput performance

	Train RMSE	Test RMSE
Daily average throughput for cache misses	19.72	6.93
Daily average throughput for cache hits	34.94	106.86
Daily cumulative total throughput for cache misses	469618.67	417960.60
Daily cumulative total throughput for cache hits	579464.01	1652297.00

Hourly LSTM model results for data access throughput performance

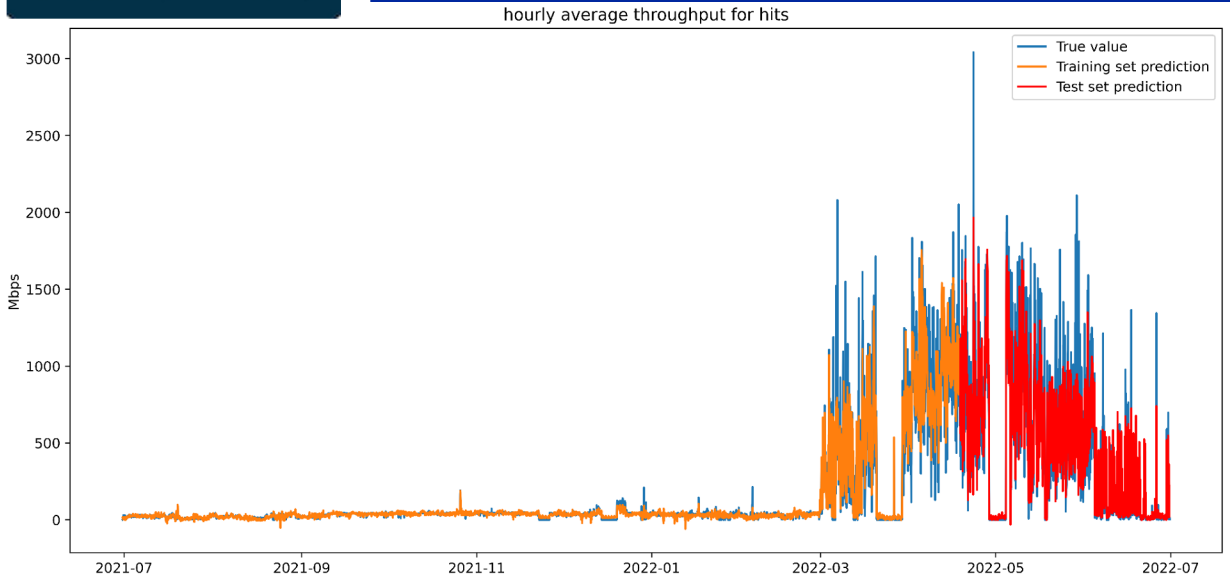

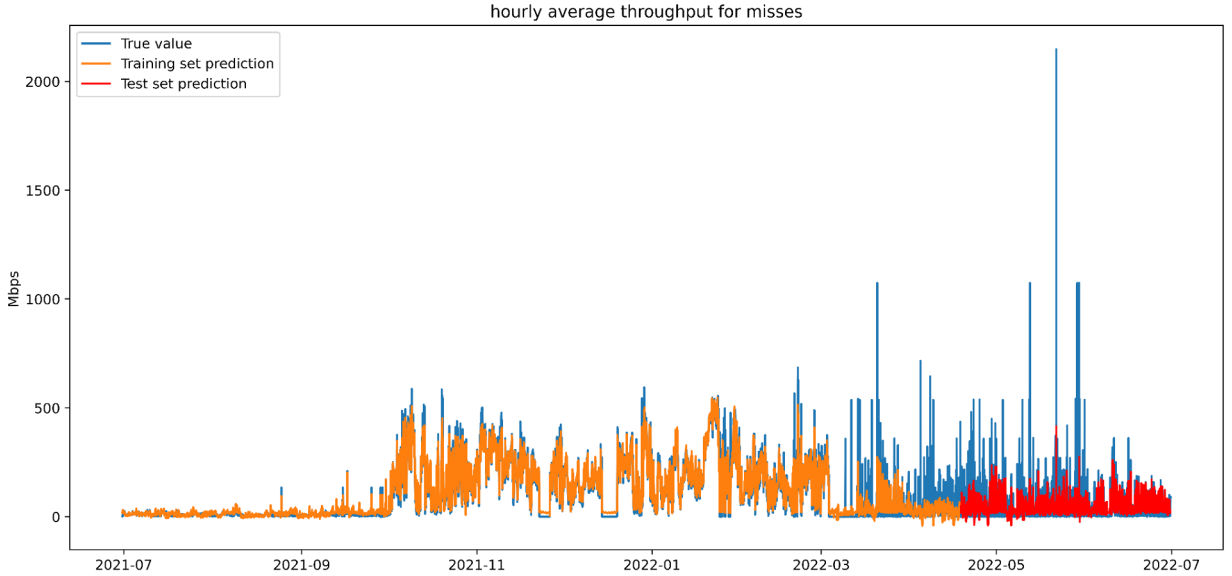

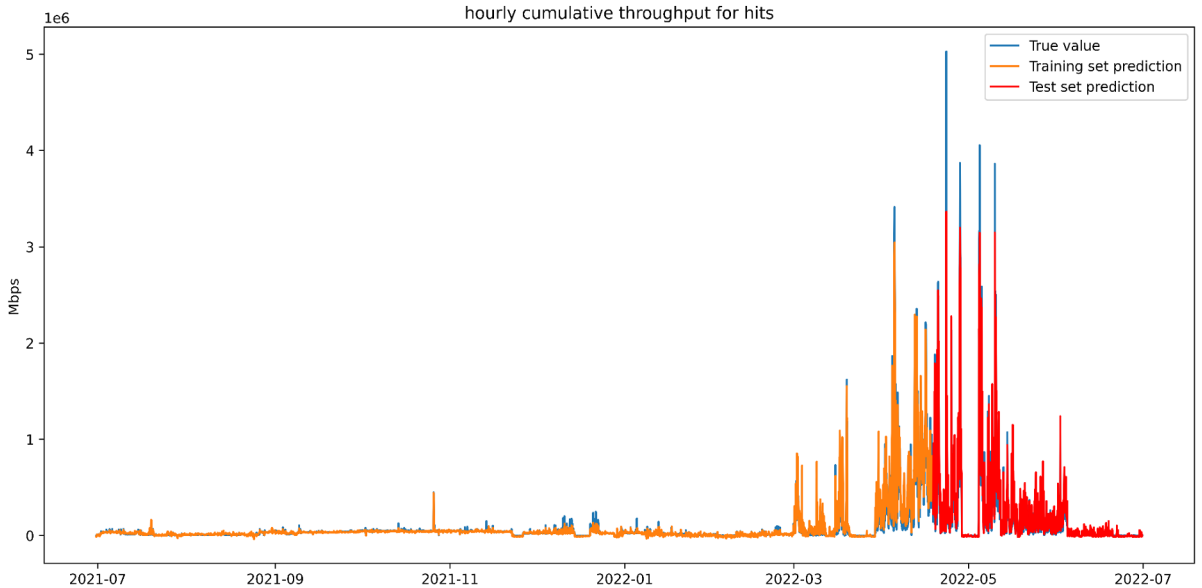

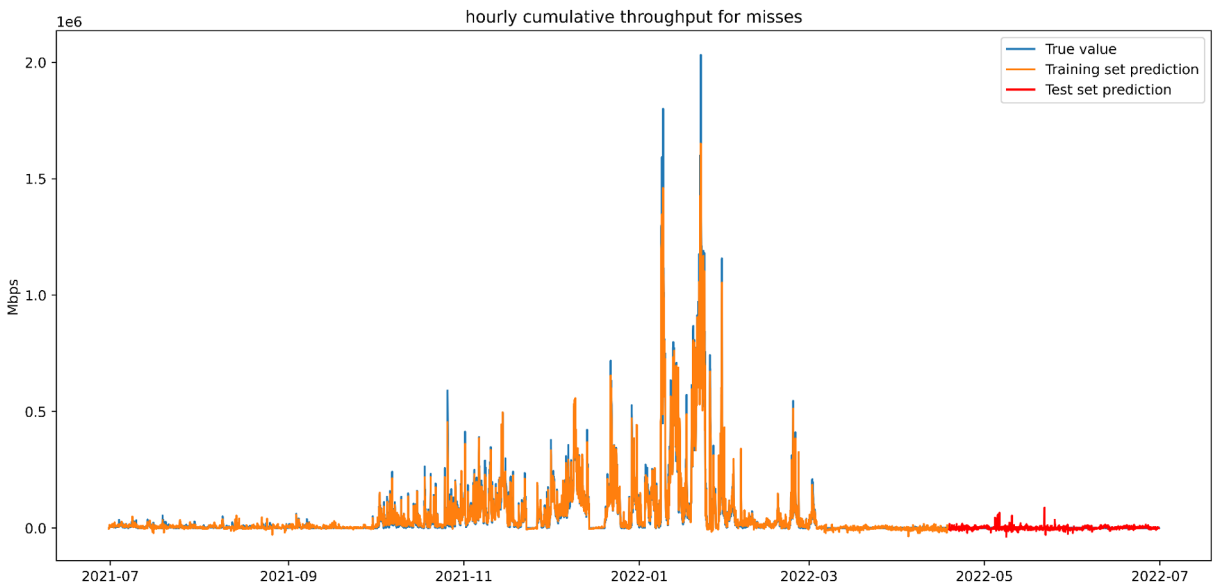

RMSE of Hourly LSTM model result for data access throughput performance

	Train RMSE	Test RMSE
Hourly average throughput for cache misses	33.26	78.87
Hourly average throughput for cache hits	62.34	136.85
Hourly cumulative total throughput for cache misses	15461.14	6652.55
Hourly cumulative total throughput for cache hits	33871.74	82204.01

Hourly LSTM model results with 24-hour moving average for data access throughput performance

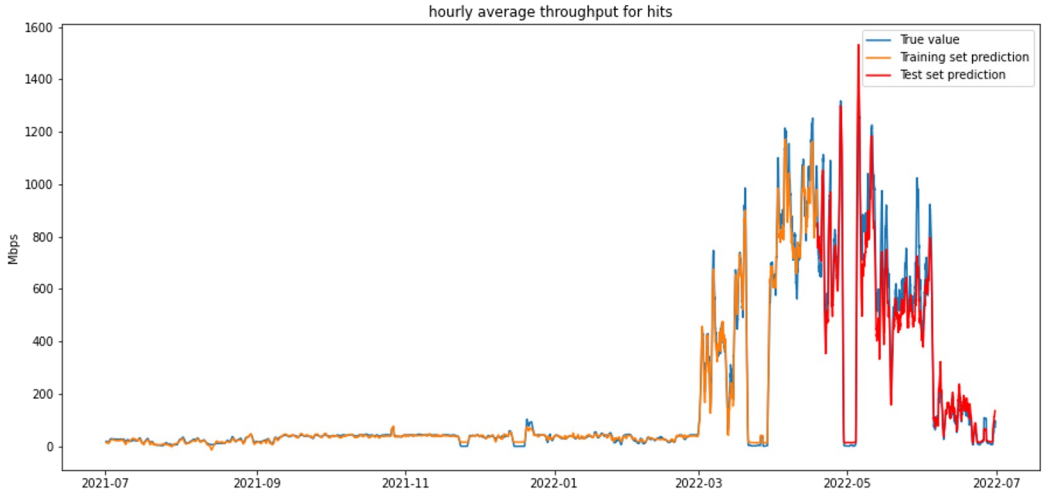

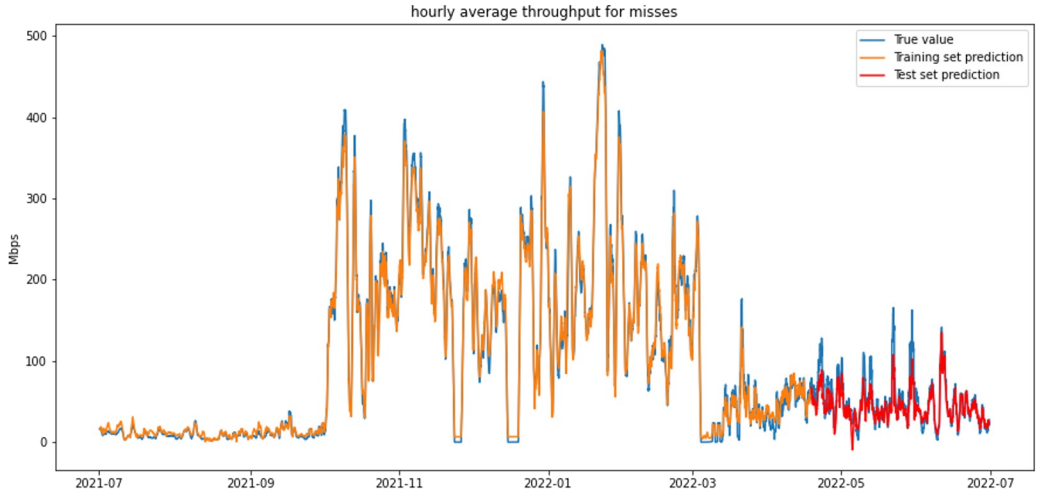

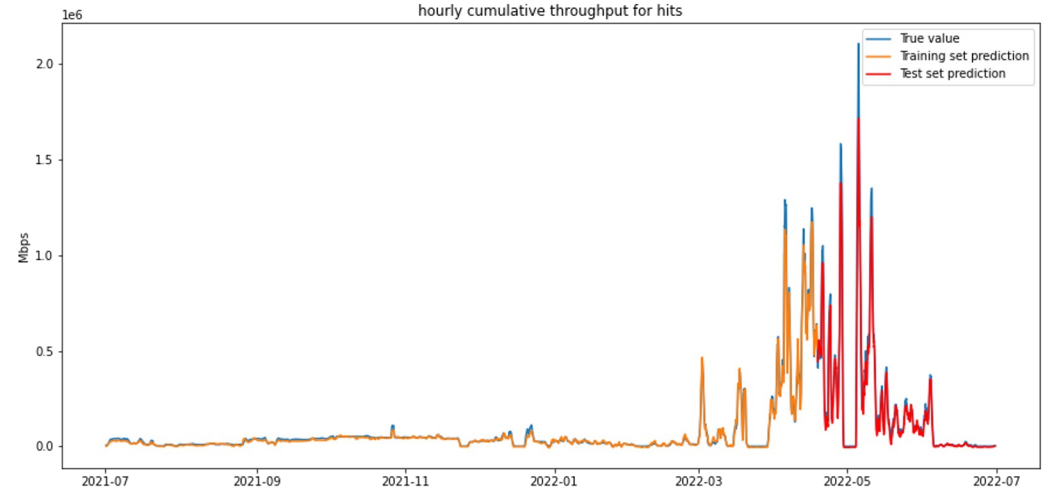

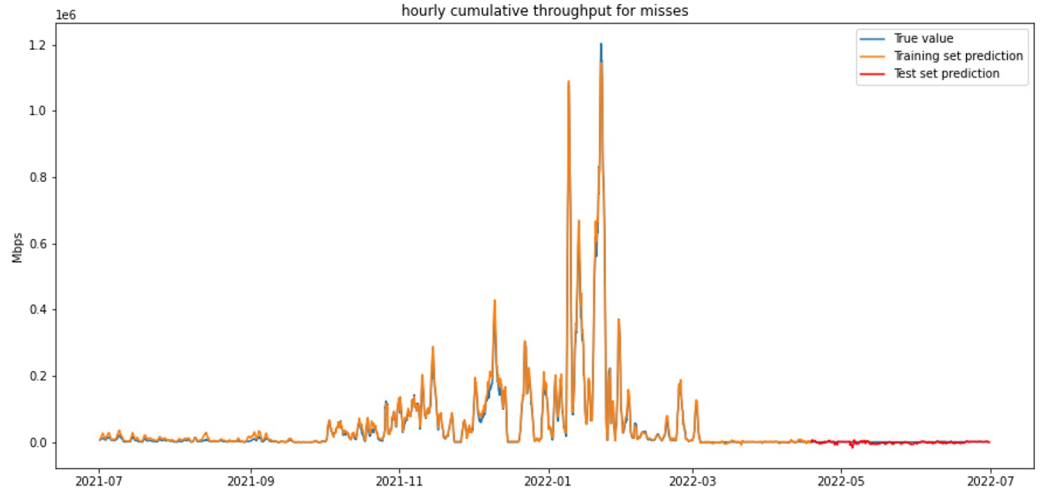

RMSE of Hourly LSTM model result with 24-hour MA for data access throughput performance

	Train RMSE	Test RMSE
Hourly average throughput for cache misses	9.76	14.10
Hourly average throughput for cache hits	18.86	81.68
Hourly cumulative total throughput for cache misses	10123.67	3338.32
Hourly cumulative total throughput for cache hits	14327.11	61091.95

Hourly LSTM model results with 168-hour moving average for data access throughput performance

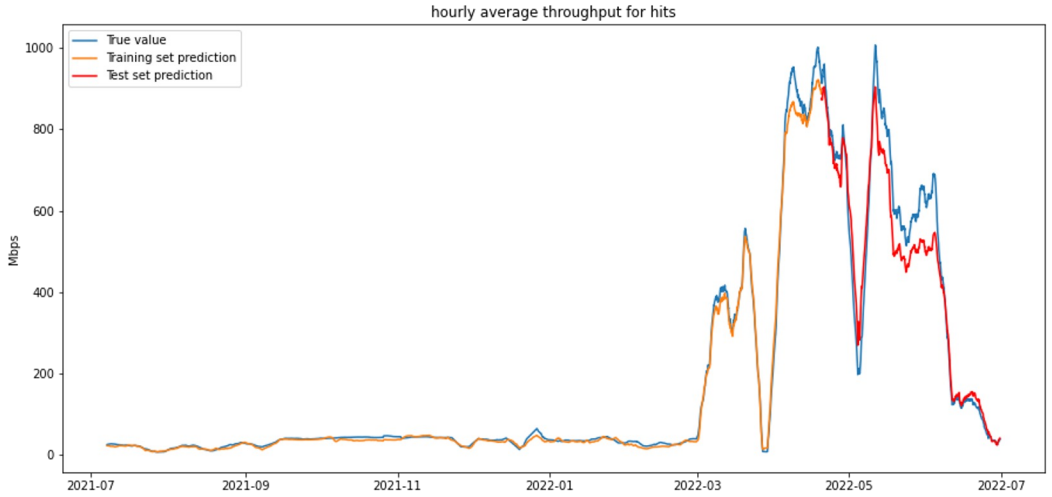

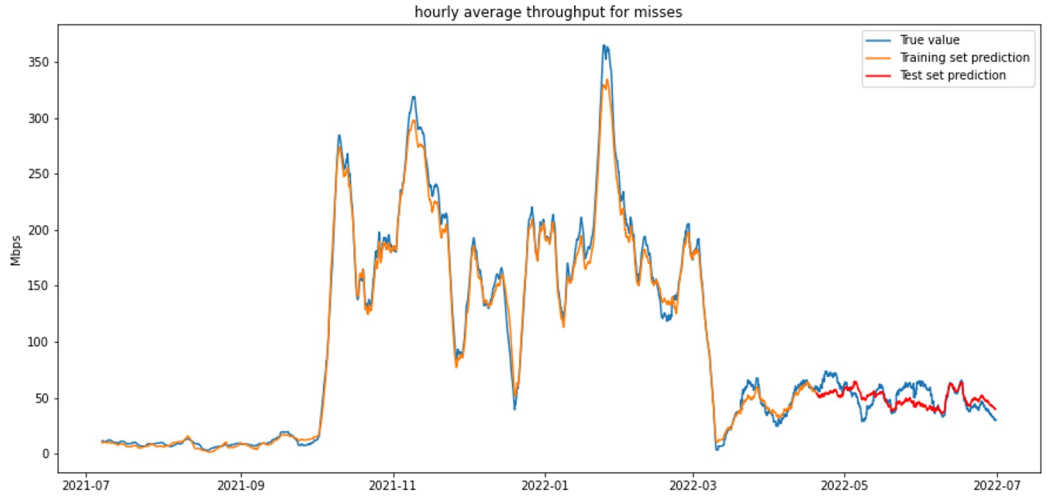

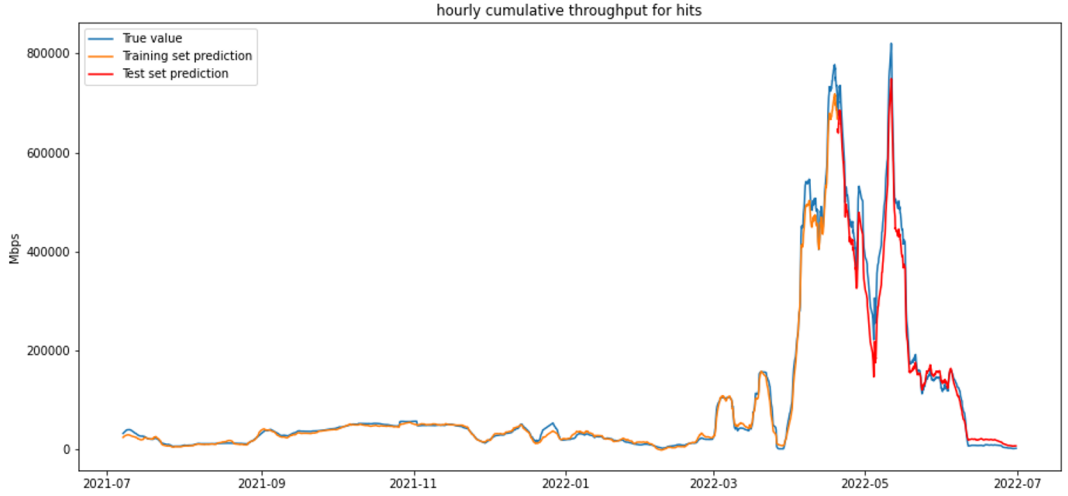

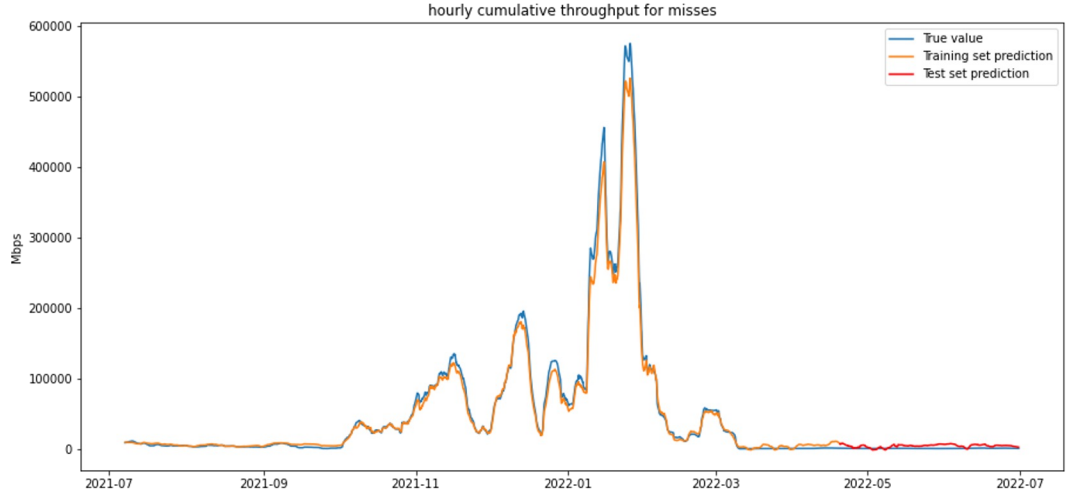

RMSE of hourly LSTM model results with 168-hour MA for data access throughput performance

	Train RMSE	Test RMSE
Hourly average throughput for cache misses	8.15	9.95
Hourly average throughput for cache hits	13.86	72.01
Hourly cumulative total throughput for cache misses	11211.42	4017.96
Hourly cumulative total throughput for cache hits	10442.30	38367.53

Summary

- General in-network regional cache could supplement the existing data repository and benefit wider user community
  - Reduced the redundant data transfers, saved network traffic volume
  - Cache utilization and network utilization
    - Reduce traffic volume by 2.35 times during normal uses
    - Reduce traffic frequency by 2.66 times
- Predictability of network resource loads and utilization
  - Cache utilization and network throughput performance can be predicted by LSTM
  - Model works better with moving average data, as there are less extreme values
  - Hourly model works better, as there are more data records
- Further studies
  - Longer term network requirements
  - Compared to data access patterns of the different regional repositories

Backup slides

Summary data accesses May - Dec 2020 (ESnet node only)

	Number of accesses	Data transfer size (GB)	Shared data access size (GB)
May 4-31, 2020	189,984	30,150.50	47,986.56
June 2020	215,452	40,835.23	55,929.47
July 2020	205,478	33,399.81	66,457.35
Aug 2020	203,806	30,819.80	68,723.19
Sep 2020	165,910	10,153.97	38,036.19
Oct 2020	306,118	22,723.93	45,614.91
Nov 2020	276	3.33	47
Dec 2020	8514	1236.81	4523
Total (May-Oct)	1,286,748	168,083.27	322,747.67
Daily average	9,674.79	1,263.8	2,426.67

Summary data accesses Jan-May 2021

	Number of accesses	Data transfer size (TiB)	Shared data size (TiB)	Percentage of shared data size
Jan 2021	1,402,696	269.62	269.62	51.42%
Feb 2021	1,078,545	279.33	173.69	38.34%
Mar 2021	1,166,506	319.77	226.57	41.47%
Apr 2021	365,068	81.85	72.04	46.81%
May 2021	757,555	216.50	186.29	46.24%
Total	4,770,370	1152.14	928.21	44.62%
Daily average	32,451.50	7.84	6.31	

CMS Data analysis

- **Data reduction**
  - **CMS reduces event size in tiered data formats**
    - **RECO->AOD (Analysis Object Data)->MiniAOD->NanoAOD**

Data Tier	Data
RAW [MB]	7.4
AOD [MB]	2.0
MiniAOD [kB]	200
NanoAOD [kB]	4

    - **E.g. 240B events/year becomes**
      ~0.5 Exabytes/year of AOD (2.0MB x $2.4e^{11}$ = ~$5e^{11}$ MB)
      ~50 Petabytes/year of MiniAOD
      ~ 1 Petabyte/year of NanoAOD