Storage Advisors Group

July 2022

Ben Kirk,
CISL / HPCD / CSG

July 13, 2022

This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.

Agenda

1. Campaign Storage
   - Lab Access / NFS - Follow-up
   - Augmentation - Procurement Process Status
2. Storage Outlook: 1-year lookahead
3. Filesystem Usage
   - Current, History, and Access Patterns
4. Roundtable

Campaign Storage

Data Sharing:

- During April's SAG, roundtable discussion proposed we explore options for NFS-type access to Campaign Storage from Lab stakeholders
- Follow-on discussion in May clarified scope & expectations:
  - Moderate quantity of mount points (~10 servers) per lab.
  - Read-only, latest NFS protocols.

Status:

- CISL has discussed and does not see any "show stoppers."
  - Concerned about scope creep - ***read-only would be a hard requirement***.
- SAG discussion: Is NFS server failover a strong desire, or is infrequent server outage acceptable?
  - Possibly trade server-side redundancy vs. urgency of feature deployment.

Near-term Augmentation Plan:

- NSF has approved procurement of 440 additional disks for a formatted capacity increase of ~4 PB capacity
  - Allocation plan in work, must also account for GLADE2 decommissioning (*next slide*).

Overall Storage 1-year Outlook

Derecho Availability:

- +60PB ClusterStor (Lustre Filesystem)

GLADE2 Transition Activities:

- End target is that GLADE2 will be retired ~3 months after Cheyenne.
- Data must be migrated to a mix of Derecho Scratch and (augmented) Campaign Storage
- Several steps along the way:
  - +60 PB available from Derecho scratch,
  - +4PB available for Campaign Storage from new disks in procurement,
    - Will be made available for GLADE2 data migration later this year
  - +40PB(+) Spectrum Scale ***license*** capacity available upon GLADE2 retirement.
    - Presents options for Campaign Storage (or other) hardware implementation

July 2022: Scratch & Campaign Storage Filesystem Usage

Filesystem Usage

- During May's SAG, details were requested regarding Campaign Storage usage history & access patterns.
- The following data are mined from CISL's `access_report.txt` files, which are archived by date at `/glade/u/access_reports`, e.g.:

```
Directory: /gpfs/csfs1
Scan date: 2022-07-12
Total Files: 1.10 B (30.56 M compressed)
Total Data: 63.0279 PB (56.4616 PB on disk)
Avg. compression ratio: 1.25 (32.9066 PB raw / 26.3366 PB compressed)
Est. compression savings: 6.5700 PB (19.97%)

Last Accessed            Data (%)             # Files (%)
-------------------------------------------------------------
<1 Month           1.7452 PB (2.77%)         21.05 M (1.92%)
1 Month           19.1360 PB (30.36%)       301.25 M (27.46%)
6 Months          13.8449 PB (21.97%)       299.11 M (27.26%)
1 Year            20.7569 PB (32.93%)       386.62 M (35.24%)
3 Years            5.1776 PB (8.21%)         75.67 M (6.90%)
5+ Years           2.3674 PB (3.76%)         13.50 M (1.23%)
```

- Processed by `/glade/u/home/benkirk/codes/jupyter_bits/disk_usage/access_reports_history.ipynb`

Data Set Size & File Count History

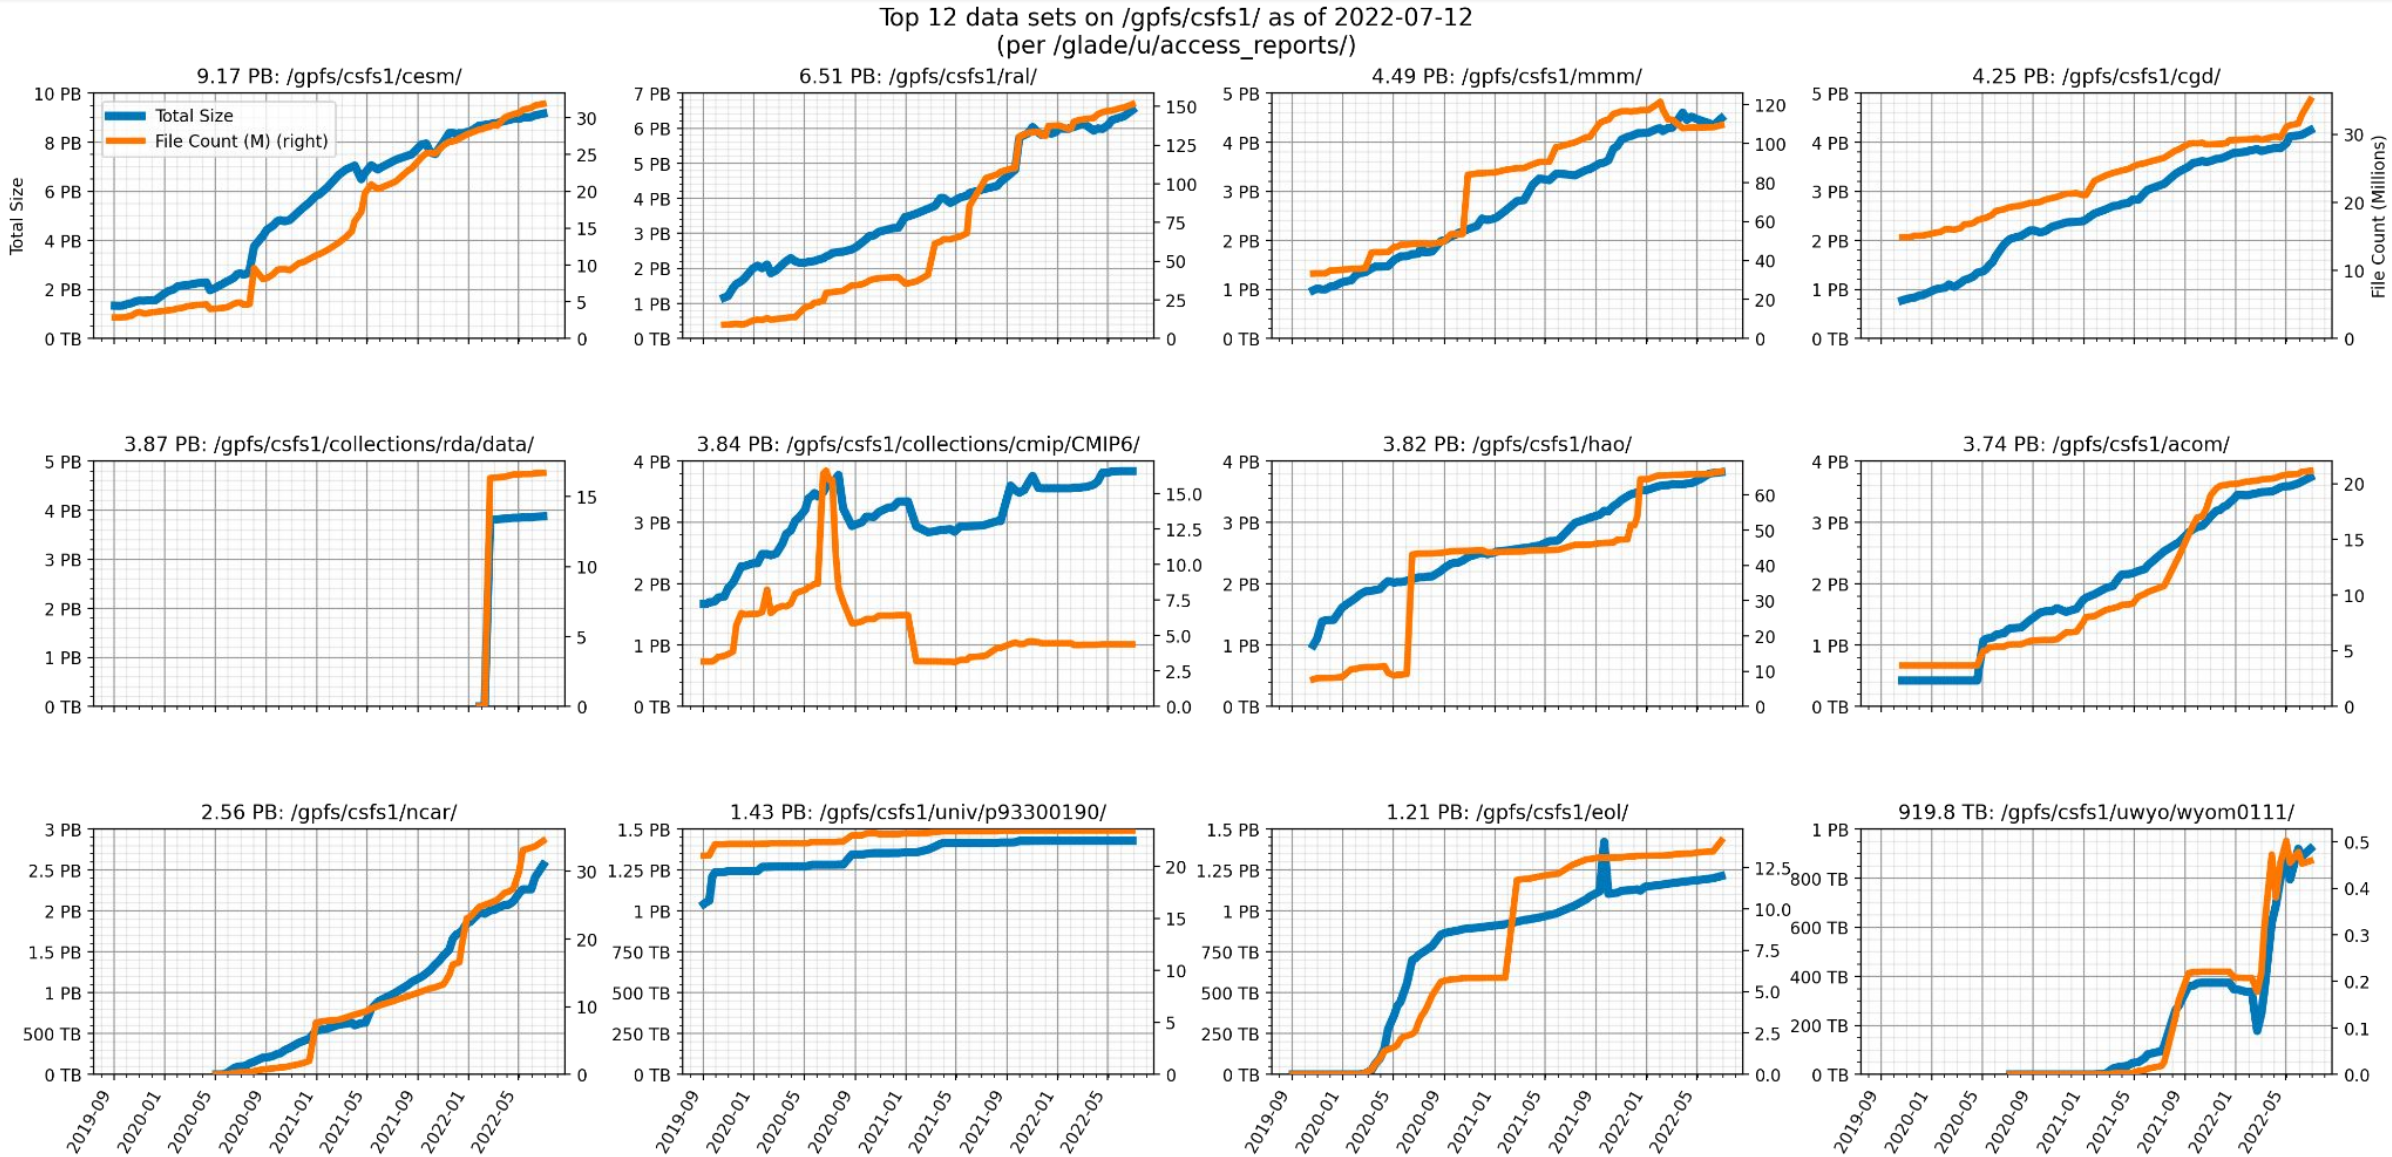

File Access Heatmap - Overview

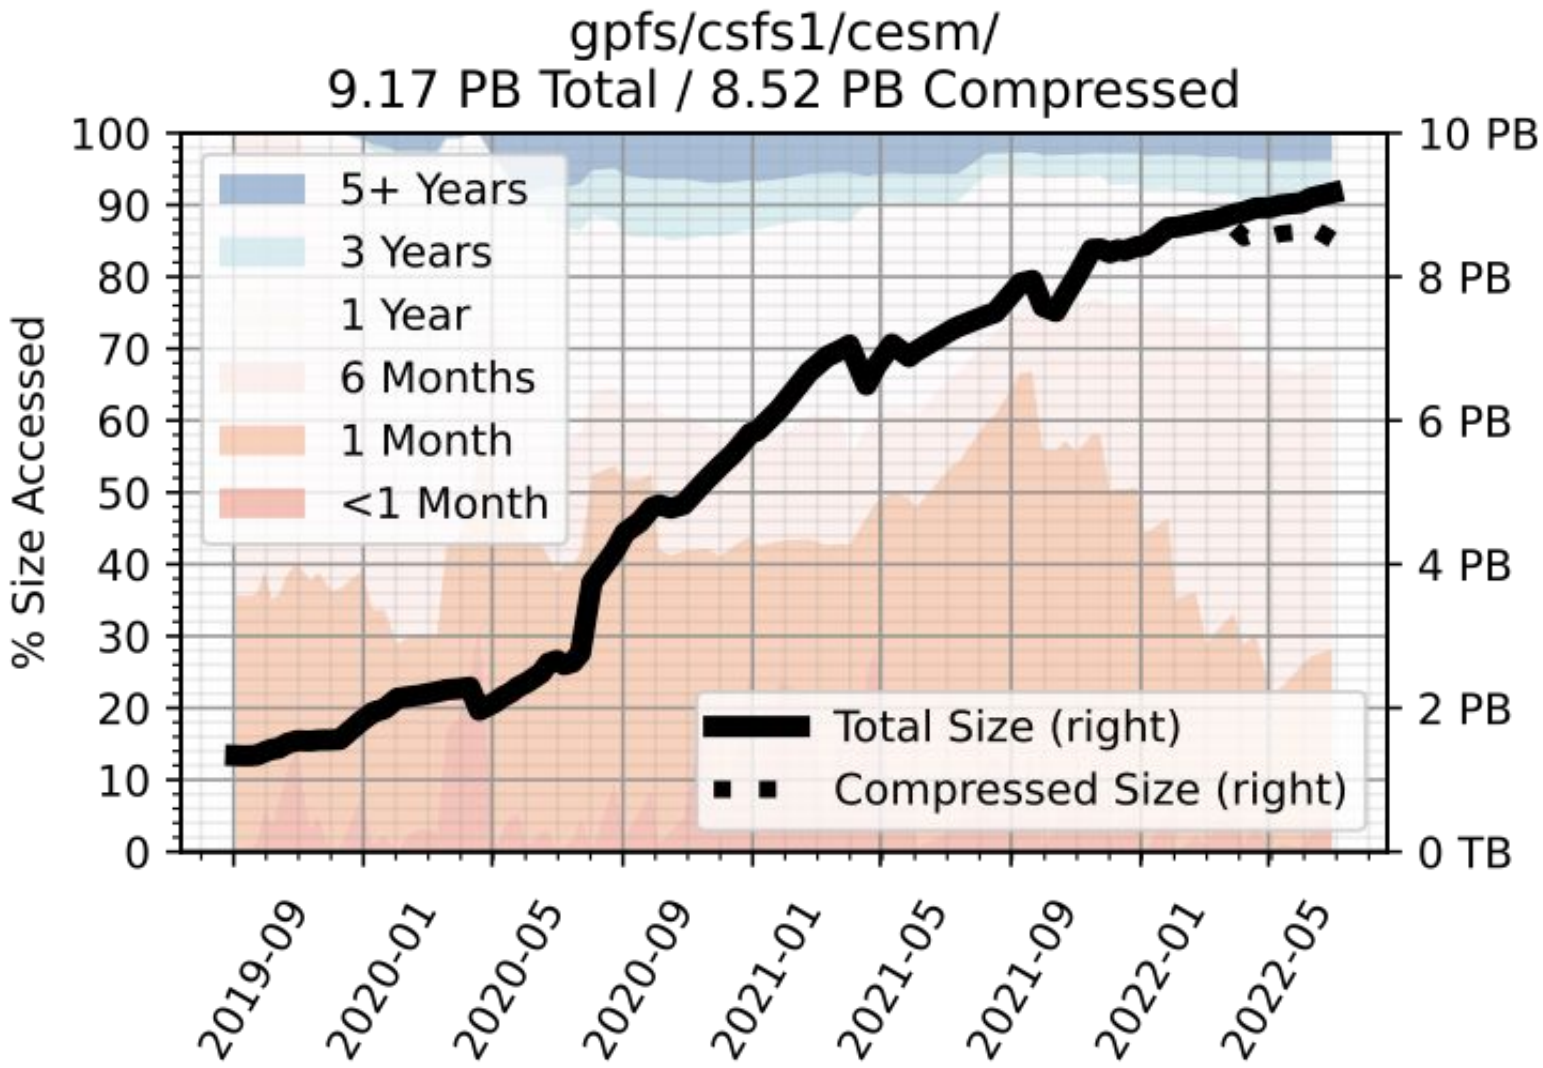

File Access Heatmap

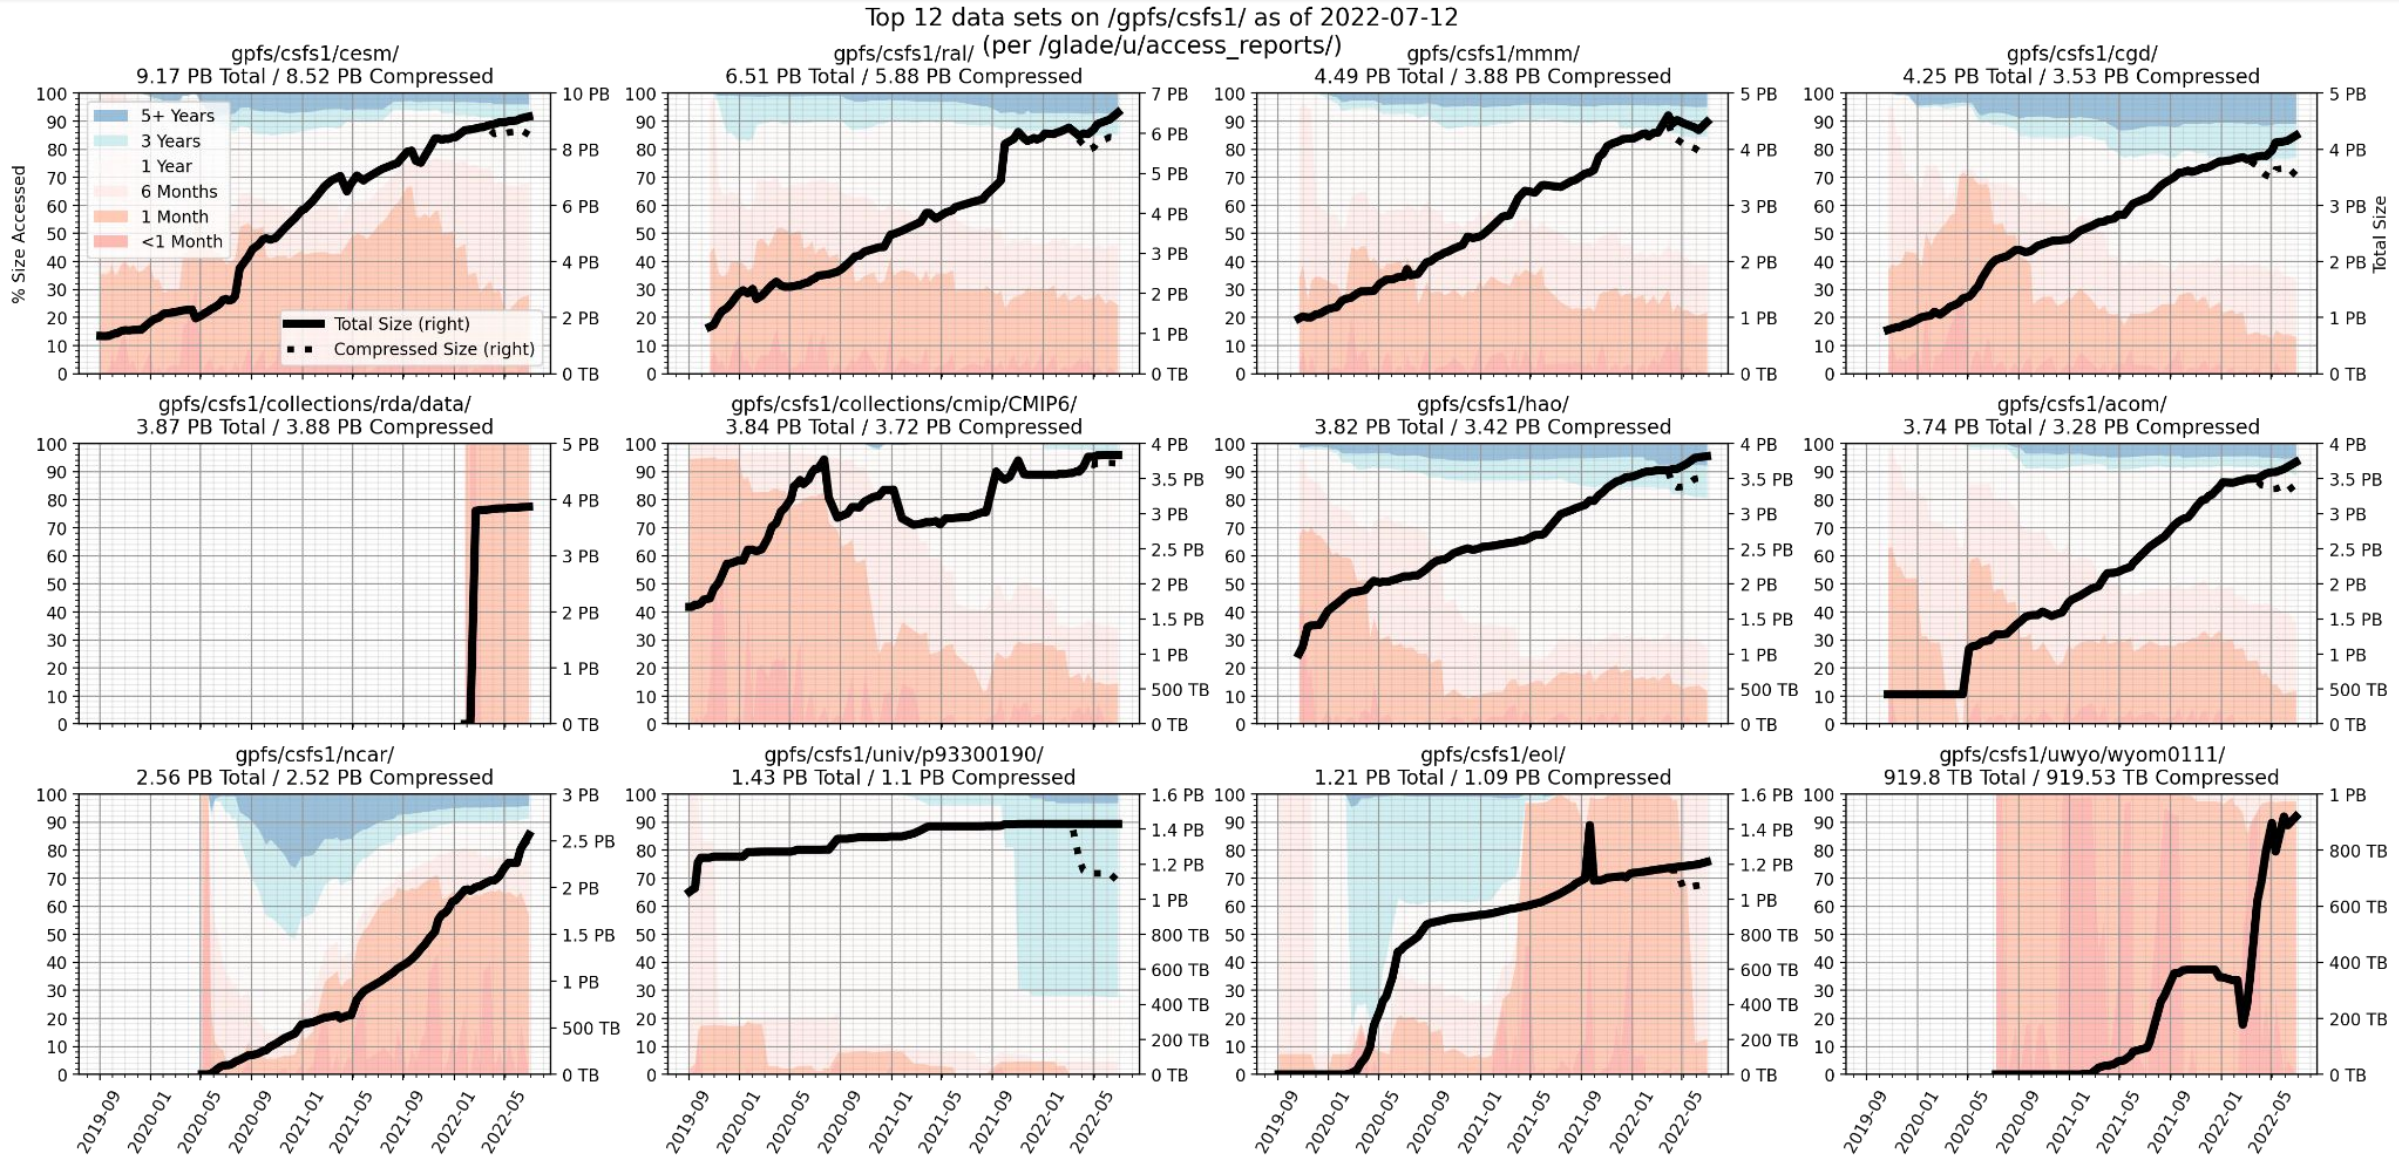

- Can we develop a policy to identify ‘cold’ Campaign Storage data at creation time, vs. years later? Thinking specifically about the ‘write once read maybe’ data use cases, or others??

- ACCLIP field campaign from July 1st through early September

Additional Questions, Issues, Concerns?