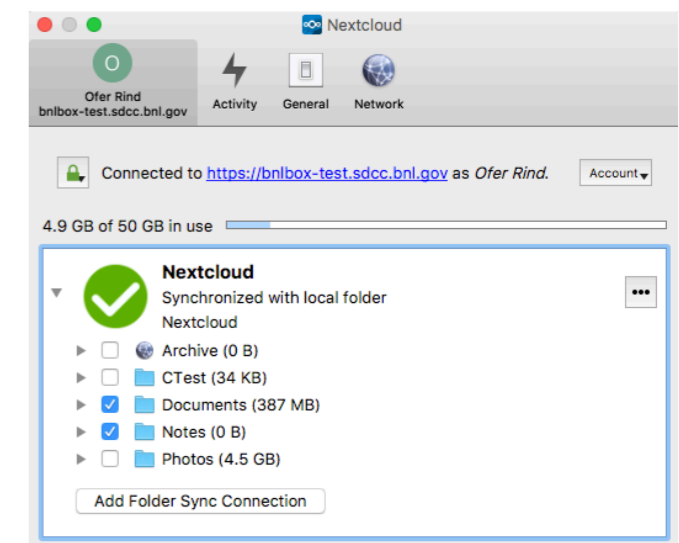
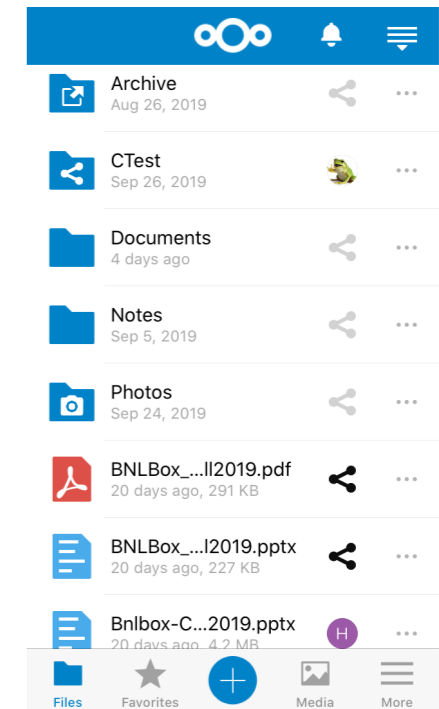



BNL BOX : Cloud Storage for Scientists

Ofer Rind
Fall HEPiX Meeting
October 12th, 2020


BNL Box is an Enterprise File Sync and Share Service (EFSS) integrated into the SDCC to provide flexible, easy-to-use, unified cloud storage for all BNL scientific users



An integrated effort by many SDCC staff supporting back-end disk and tape storage, web and db services, AAI front-end, user interface, etc.

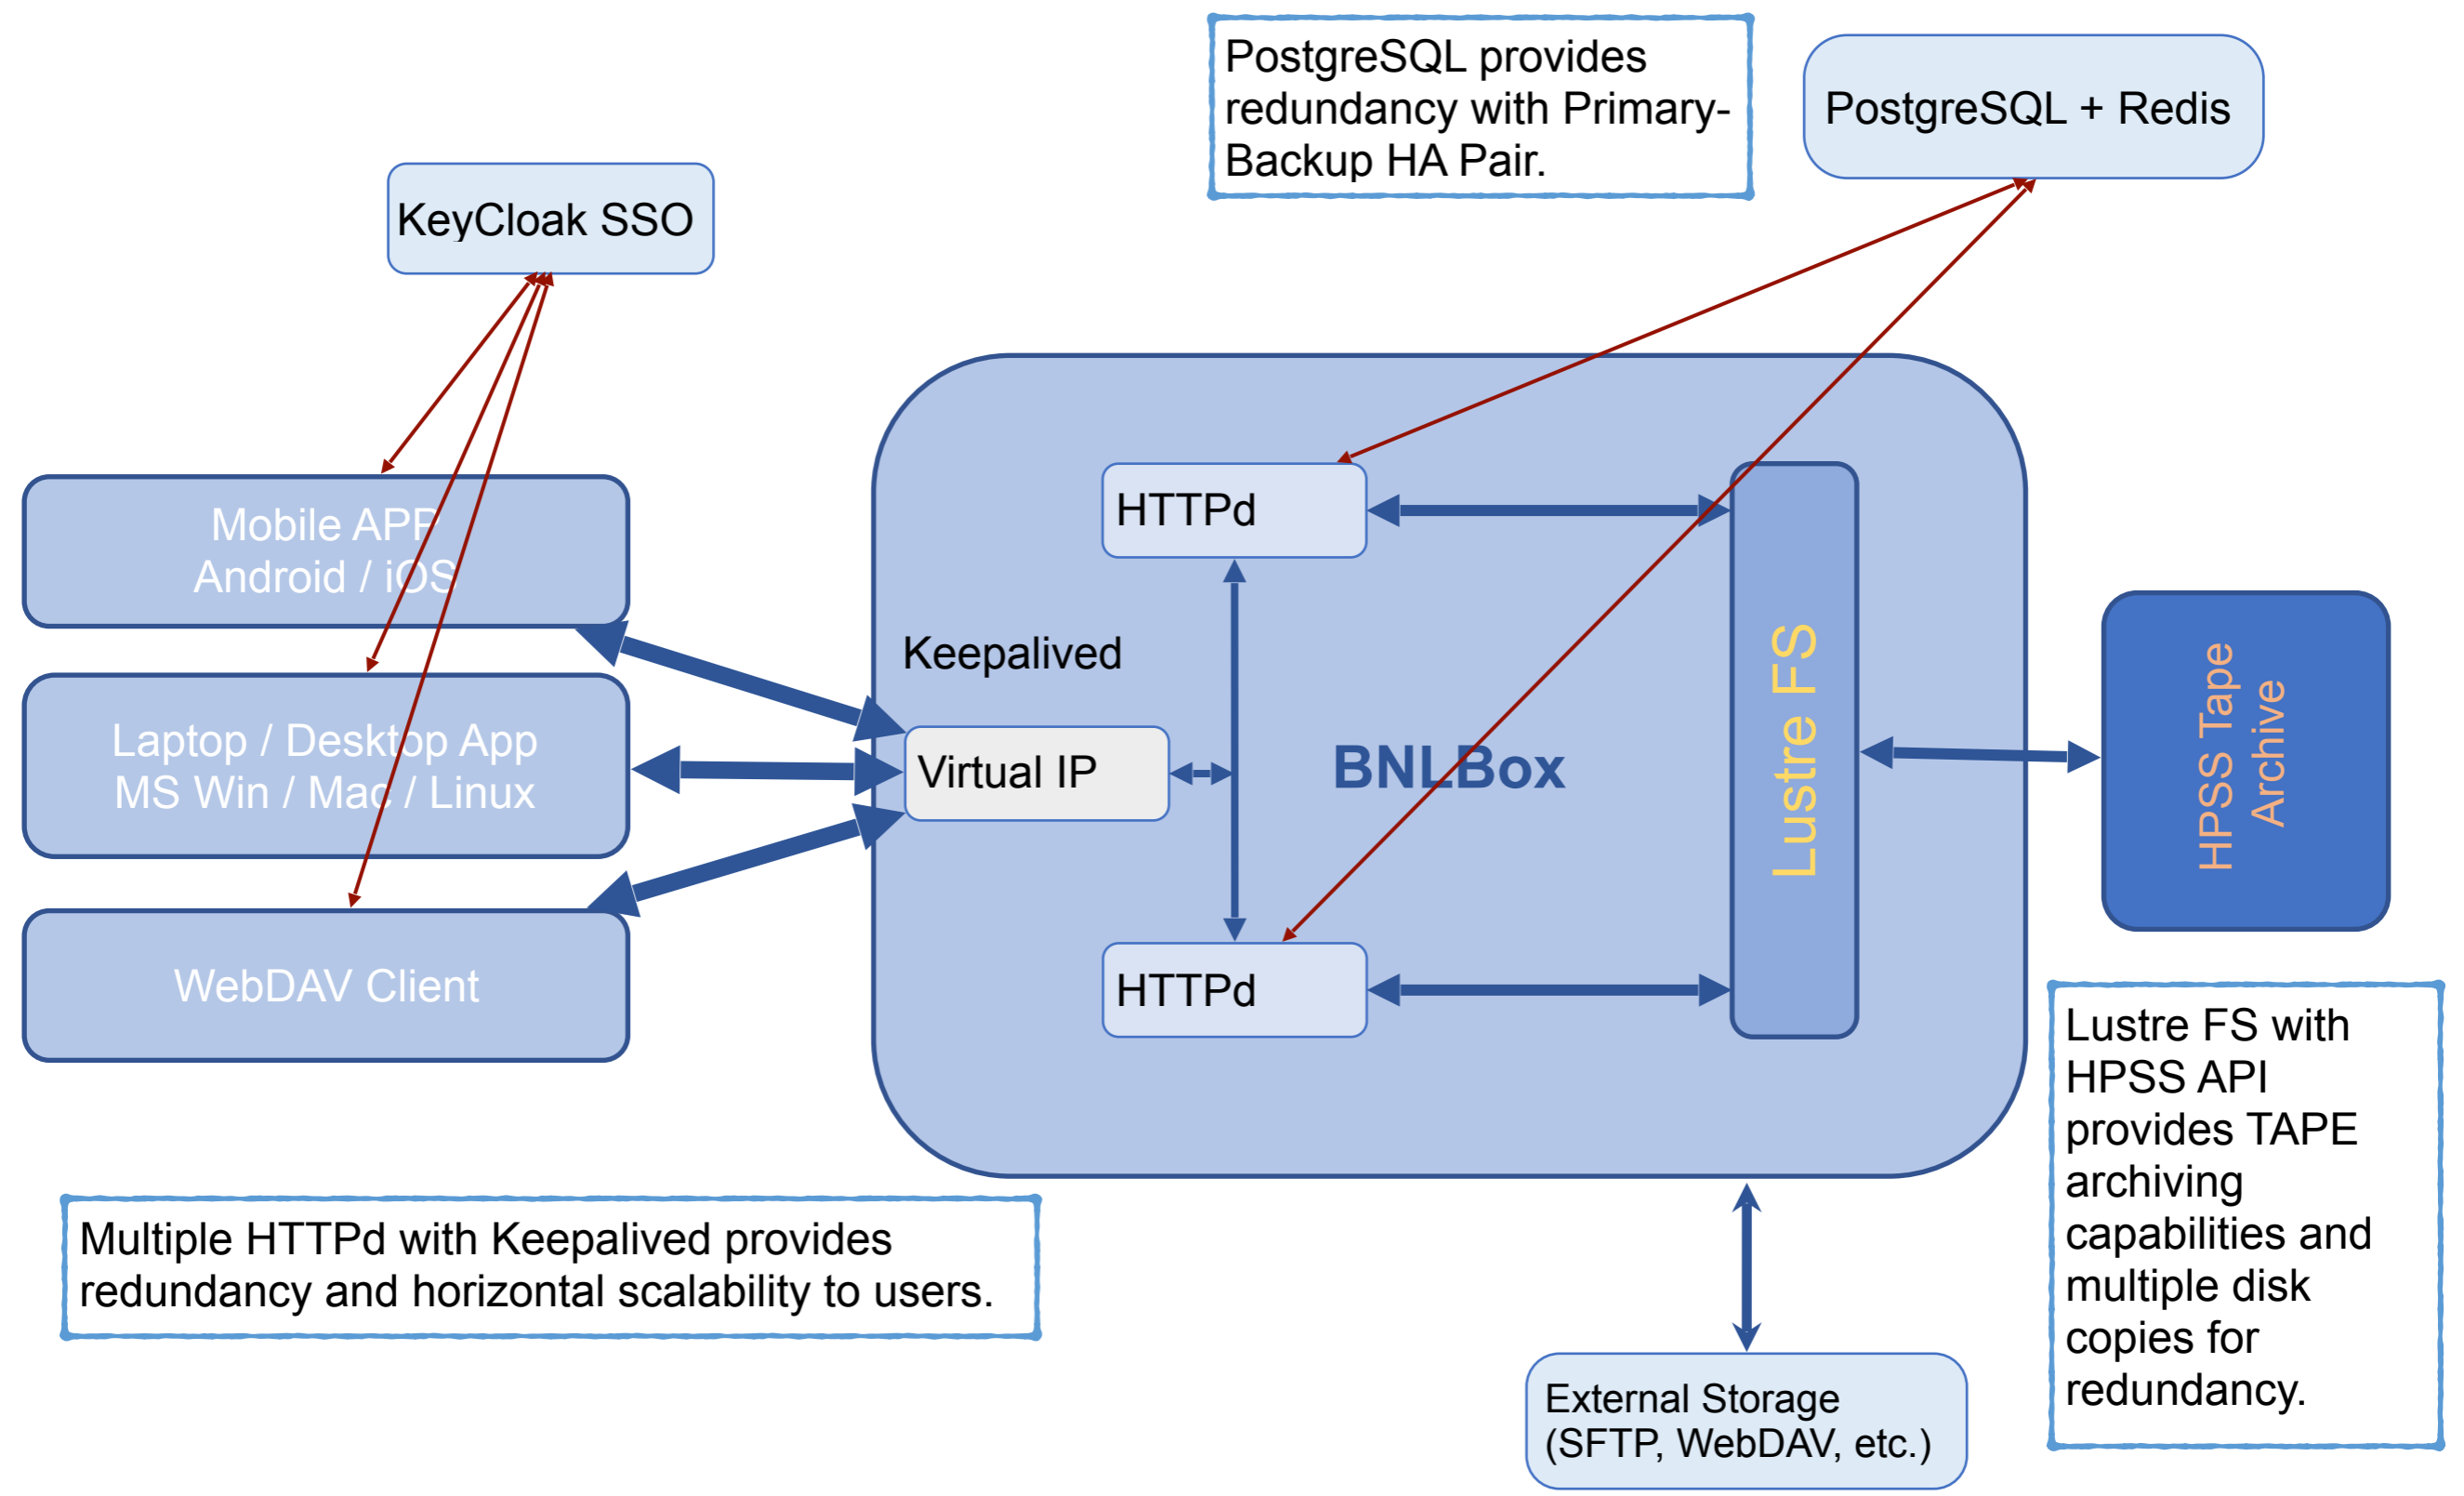
- BNL is a multi-disciplinary lab with large and small groups having diverse storage needs
- Some were already struggling with the limitations and costs of ad-hoc solutions for sharing data with collaborators, both within and outside the lab
 - Multiple access points from local user accounts, batch system, analysis portal, etc.
 - More flexible mechanisms for managing group shares
- Others were looking for easy, automated solutions for moving data to SDCC (from beamlines and other instruments)
 - Potential for archiving data to tape

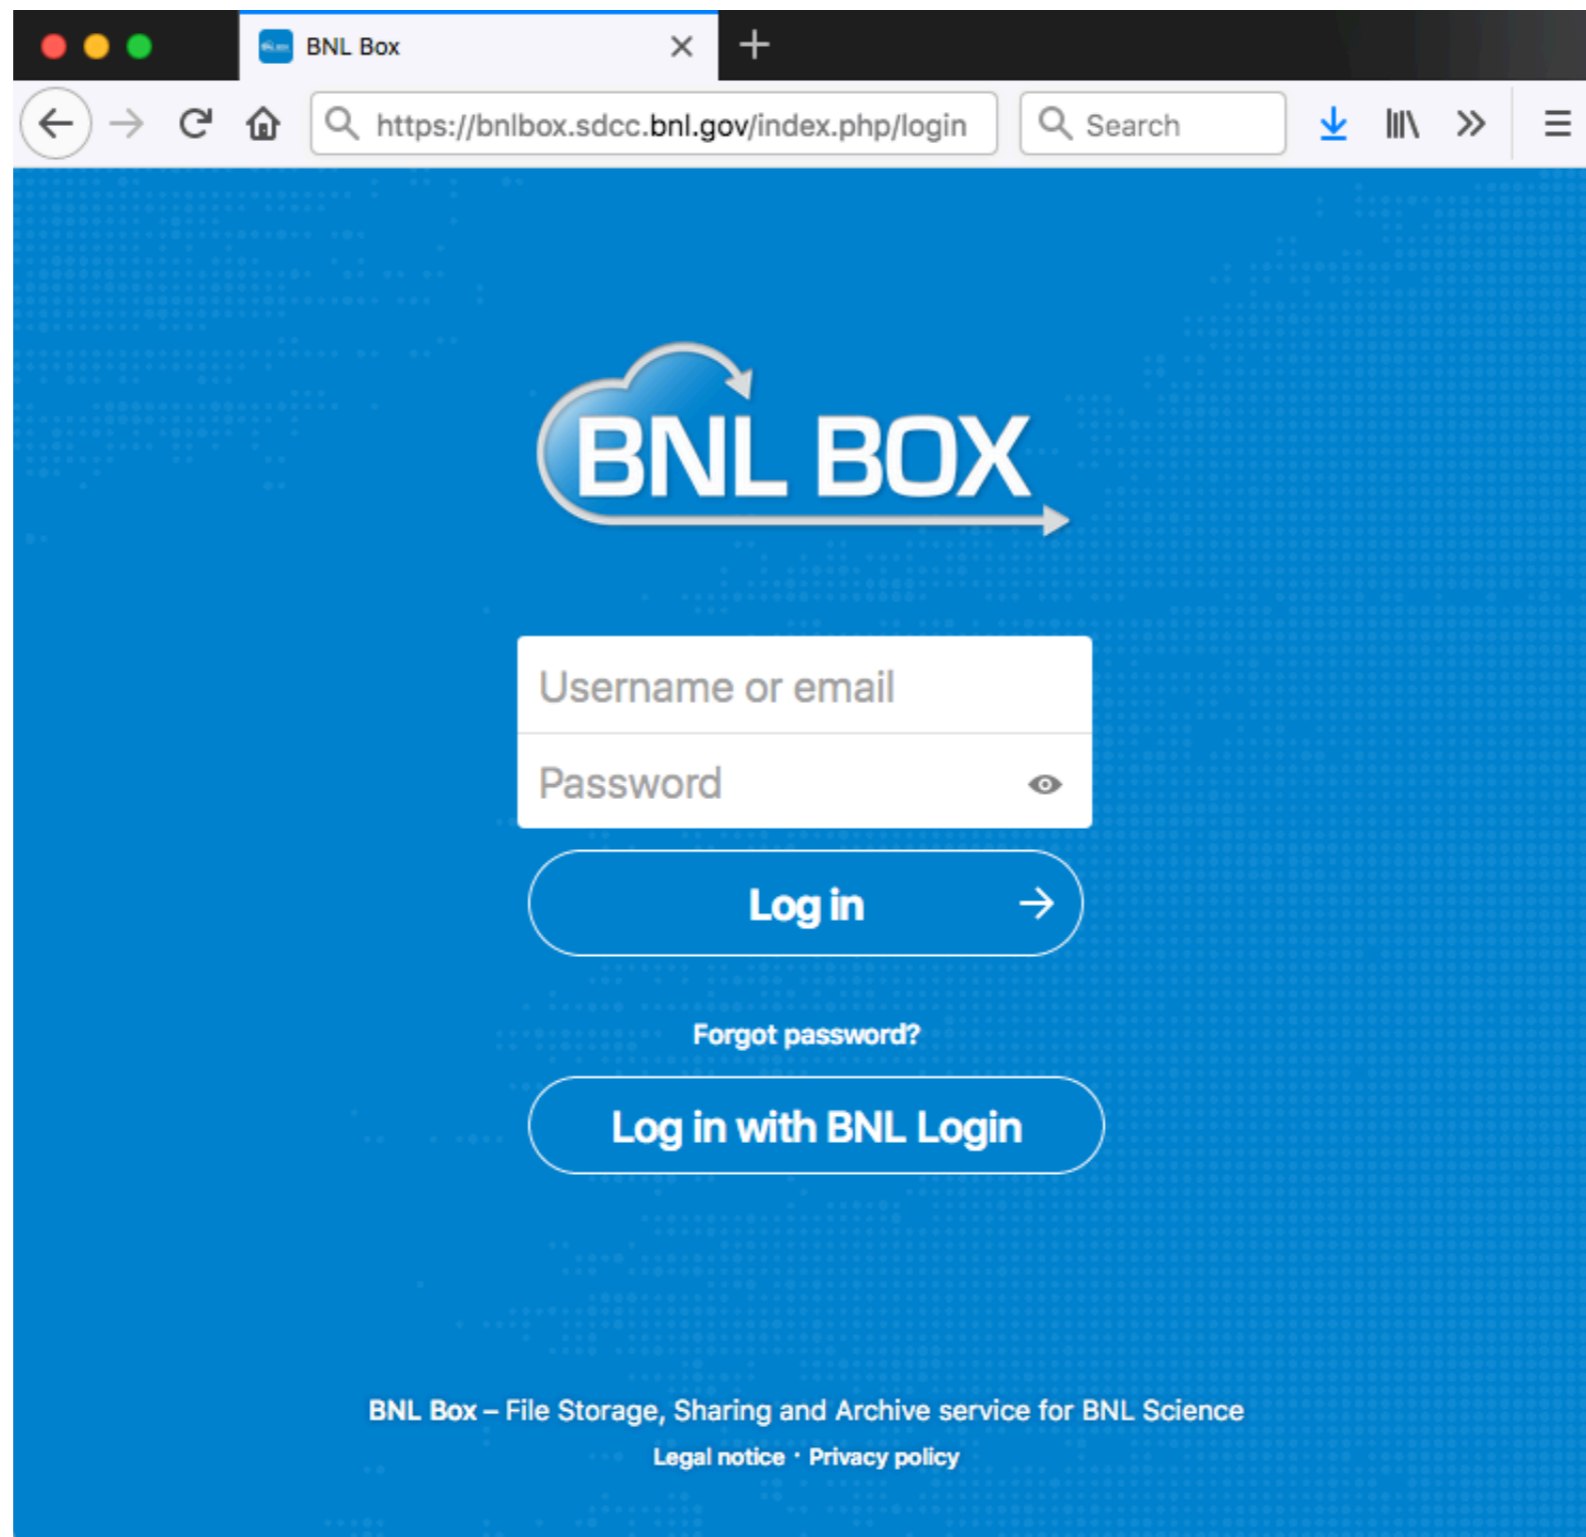
- **Belle-II:** Storing DIRAC production logs and sharing with production team at KEK
 - **EIC:** Making TBs of pre-generated Monte Carlo Data available to collaborators for download via webdav
 - **CFN:** Syncing and sharing images from AFM, SEM and other devices
- **Sync vs. Copy**
Sync is customizable but requires dedicated space on client side - not appropriate for large datasets



The first-ever direct image of a chemical structure —specifically of pentacene, an aromatic hydrocarbon made of five rings fused in a linear arrangement. *Credit: IBM.*


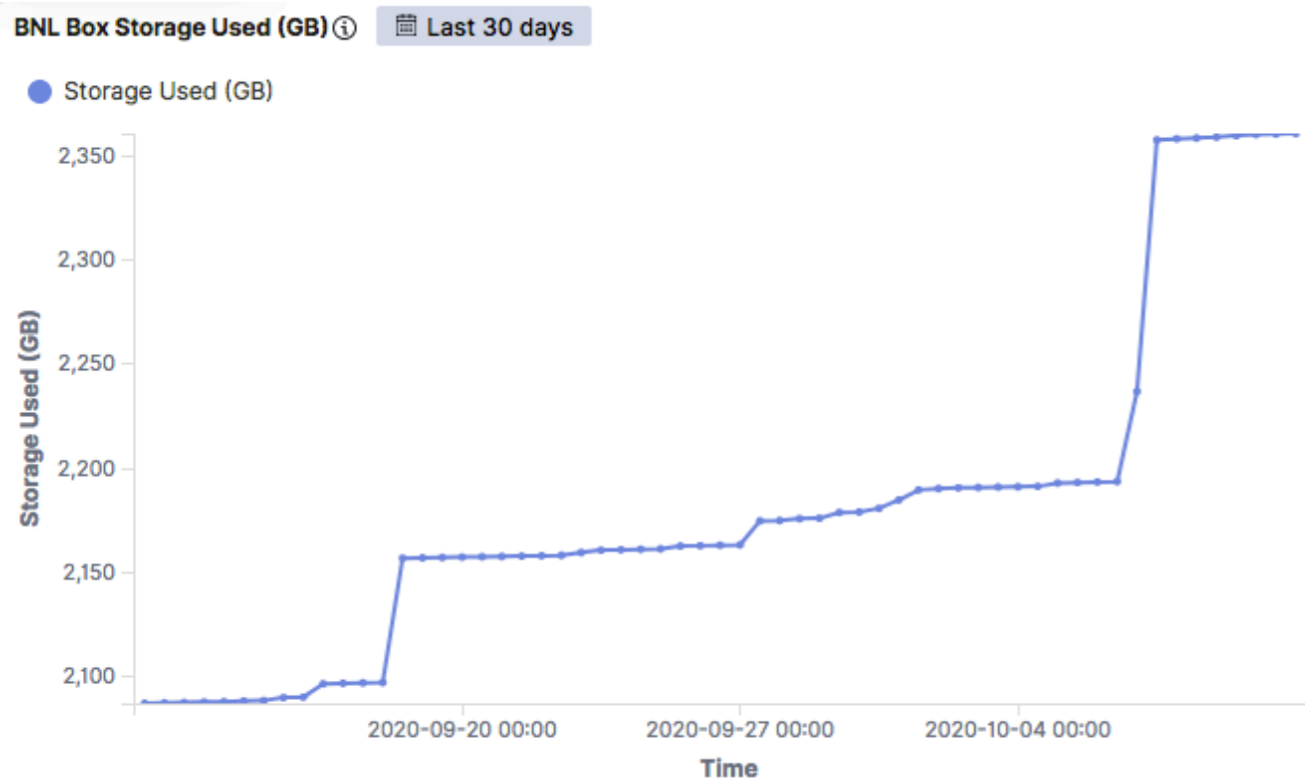
- BNLBox became a production service of the SDCC in Nov 2019, replacing the OwnCloud/Ceph prototype setup that had been running for the previous two years (all users migrated from old system by Feb 2020)
 - Based on Nextcloud version 17. 100% free open source software (forked from OwnCloud in 2016).
 - PostgreSQL+Redis backend with a HA Primary/Backup pair for redundancy, and fast NVMe flash storage for performance.
 - Lustre file system, with multiple copies for redundancy, an HPSS API for archival storage and TSM backup. **Currently 1PB of usable storage**
 - Apache server front-end pair using KeepAlived for redundancy and load-balancing with shared config directory (requires flock mount option for Lustre)
 - 40 Gbps WAN-to-storage network capacity
 - Authentication via KeyCloak OIDC integrated with SDCC LDAP and BNL AD
- **Account with 50 GB default quota automatically available to any BNL user**



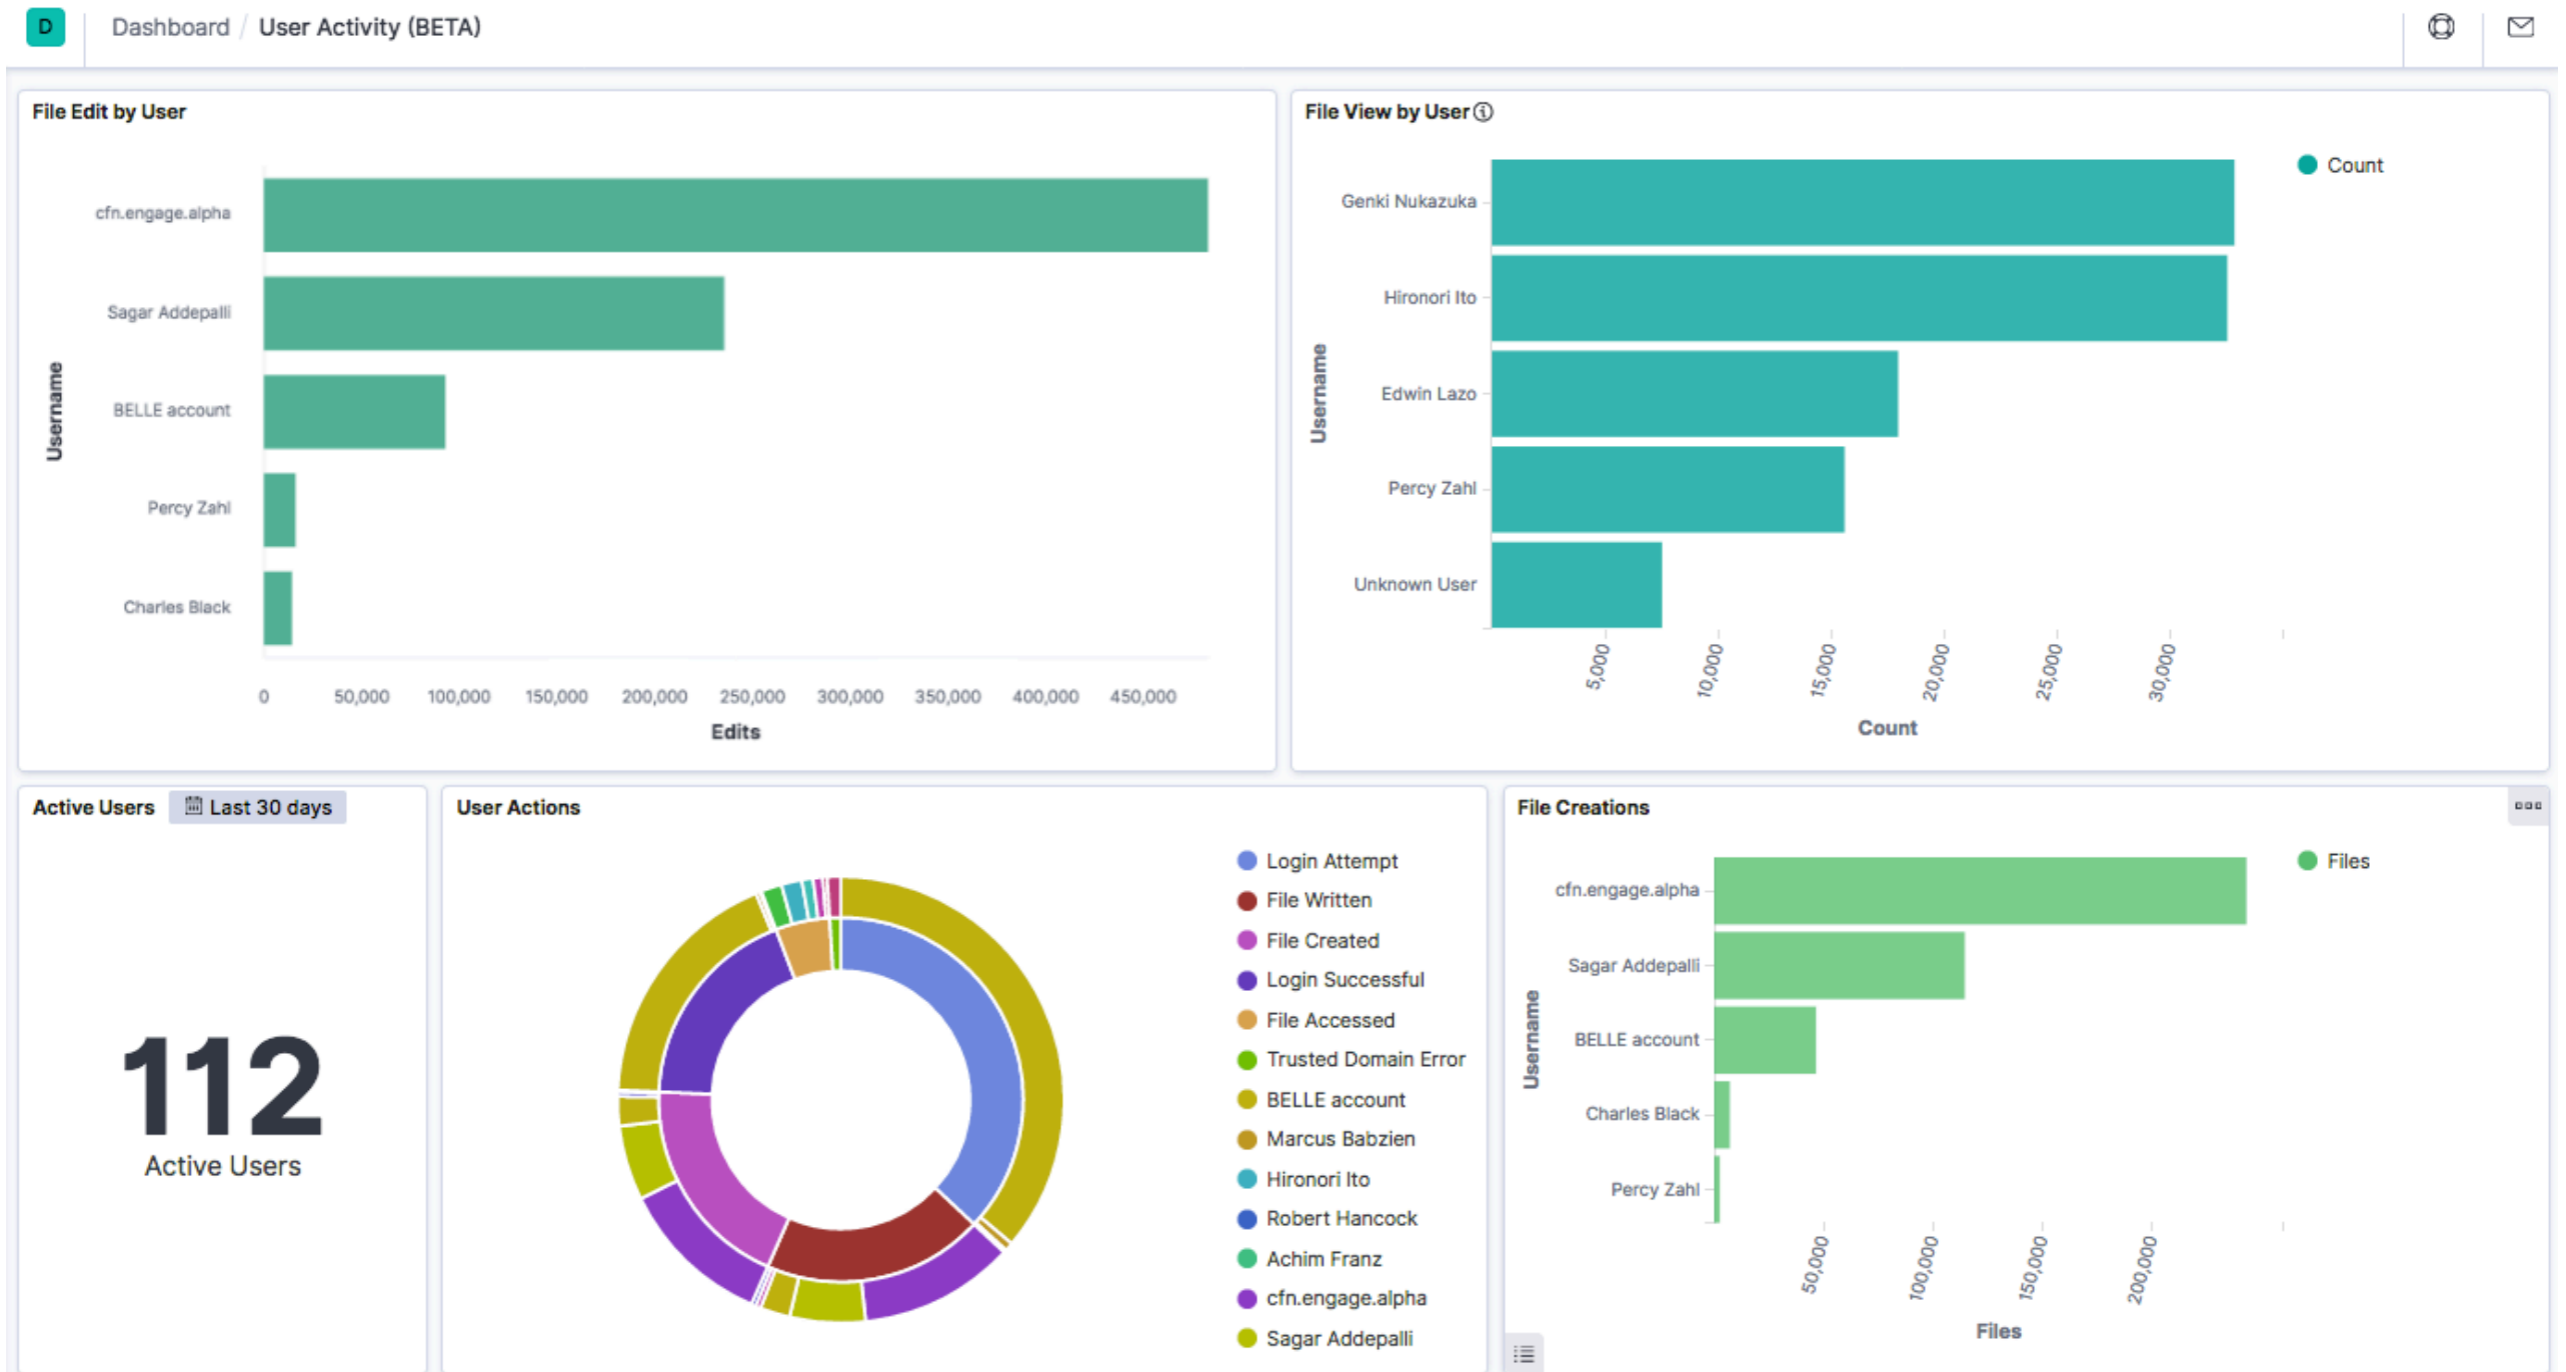


<https://bnlbox.sdcc.bnl.gov>

- Currently 213 users, plus guest accounts
- ~2 TB of user data comprising ~4M files (82% of space used by top 5 users)
- Standard Ganglia and Nagios system monitoring
- Instrumenting/testing ELK dashboards using metricbeats, filebeats, custom DB queries (B. Weindorf, M. Snyder)



User Activity Dashboard - Last 30 Days



Also monitoring activity by file (most downloaded, edited, etc.)

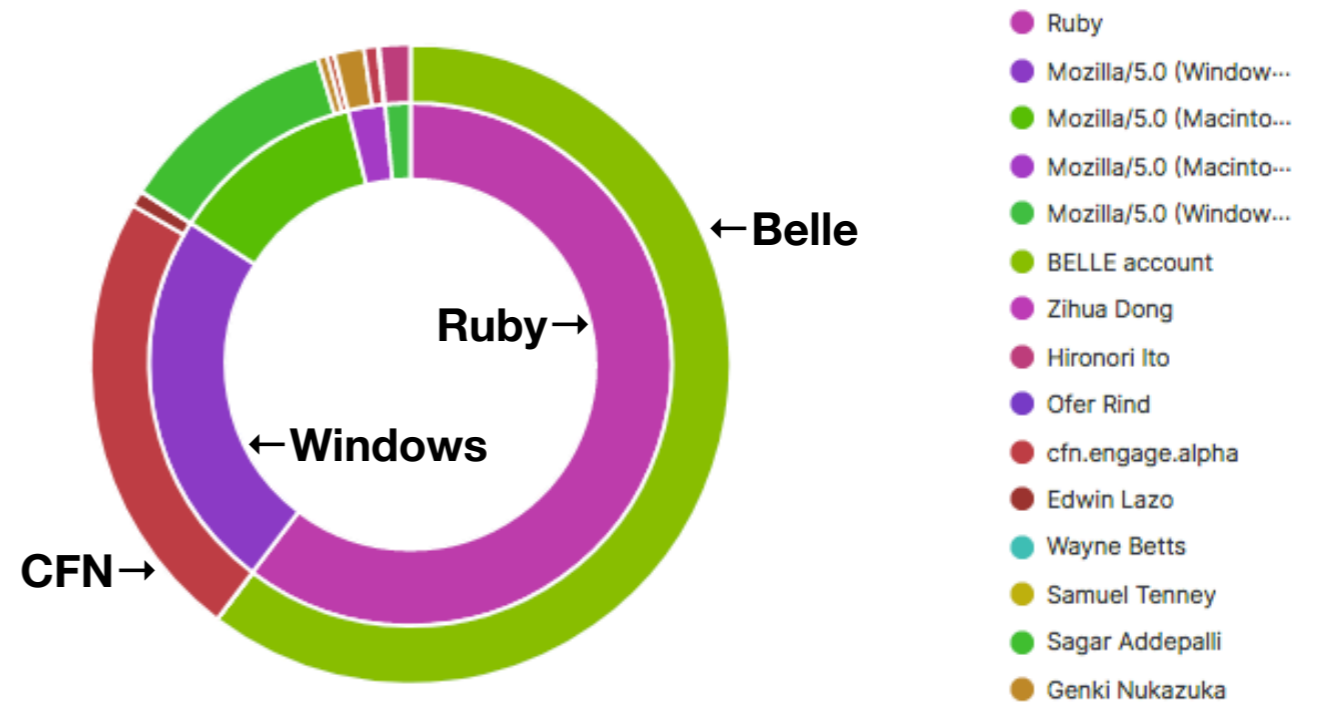
Visualizing top user activity broken down by client and type of file share creation

User Events ⓘ


Username ↕	Events ↕
BELLE account	1,271,458
cfn.engage.alpha	483,325
Sagar Addepalli	236,457
Genki Nukazuka	43,511
Hironori Ito	36,141
Percy Zahl	32,243
Unknown User	30,700
Marcus Babzien	20,682
Edwin Lazo	18,731
Charles Black	14,906

Sagar Addepalli
BELLE account
 cfn.engage.alpha
 Genki Nukazuka Hironori Ito

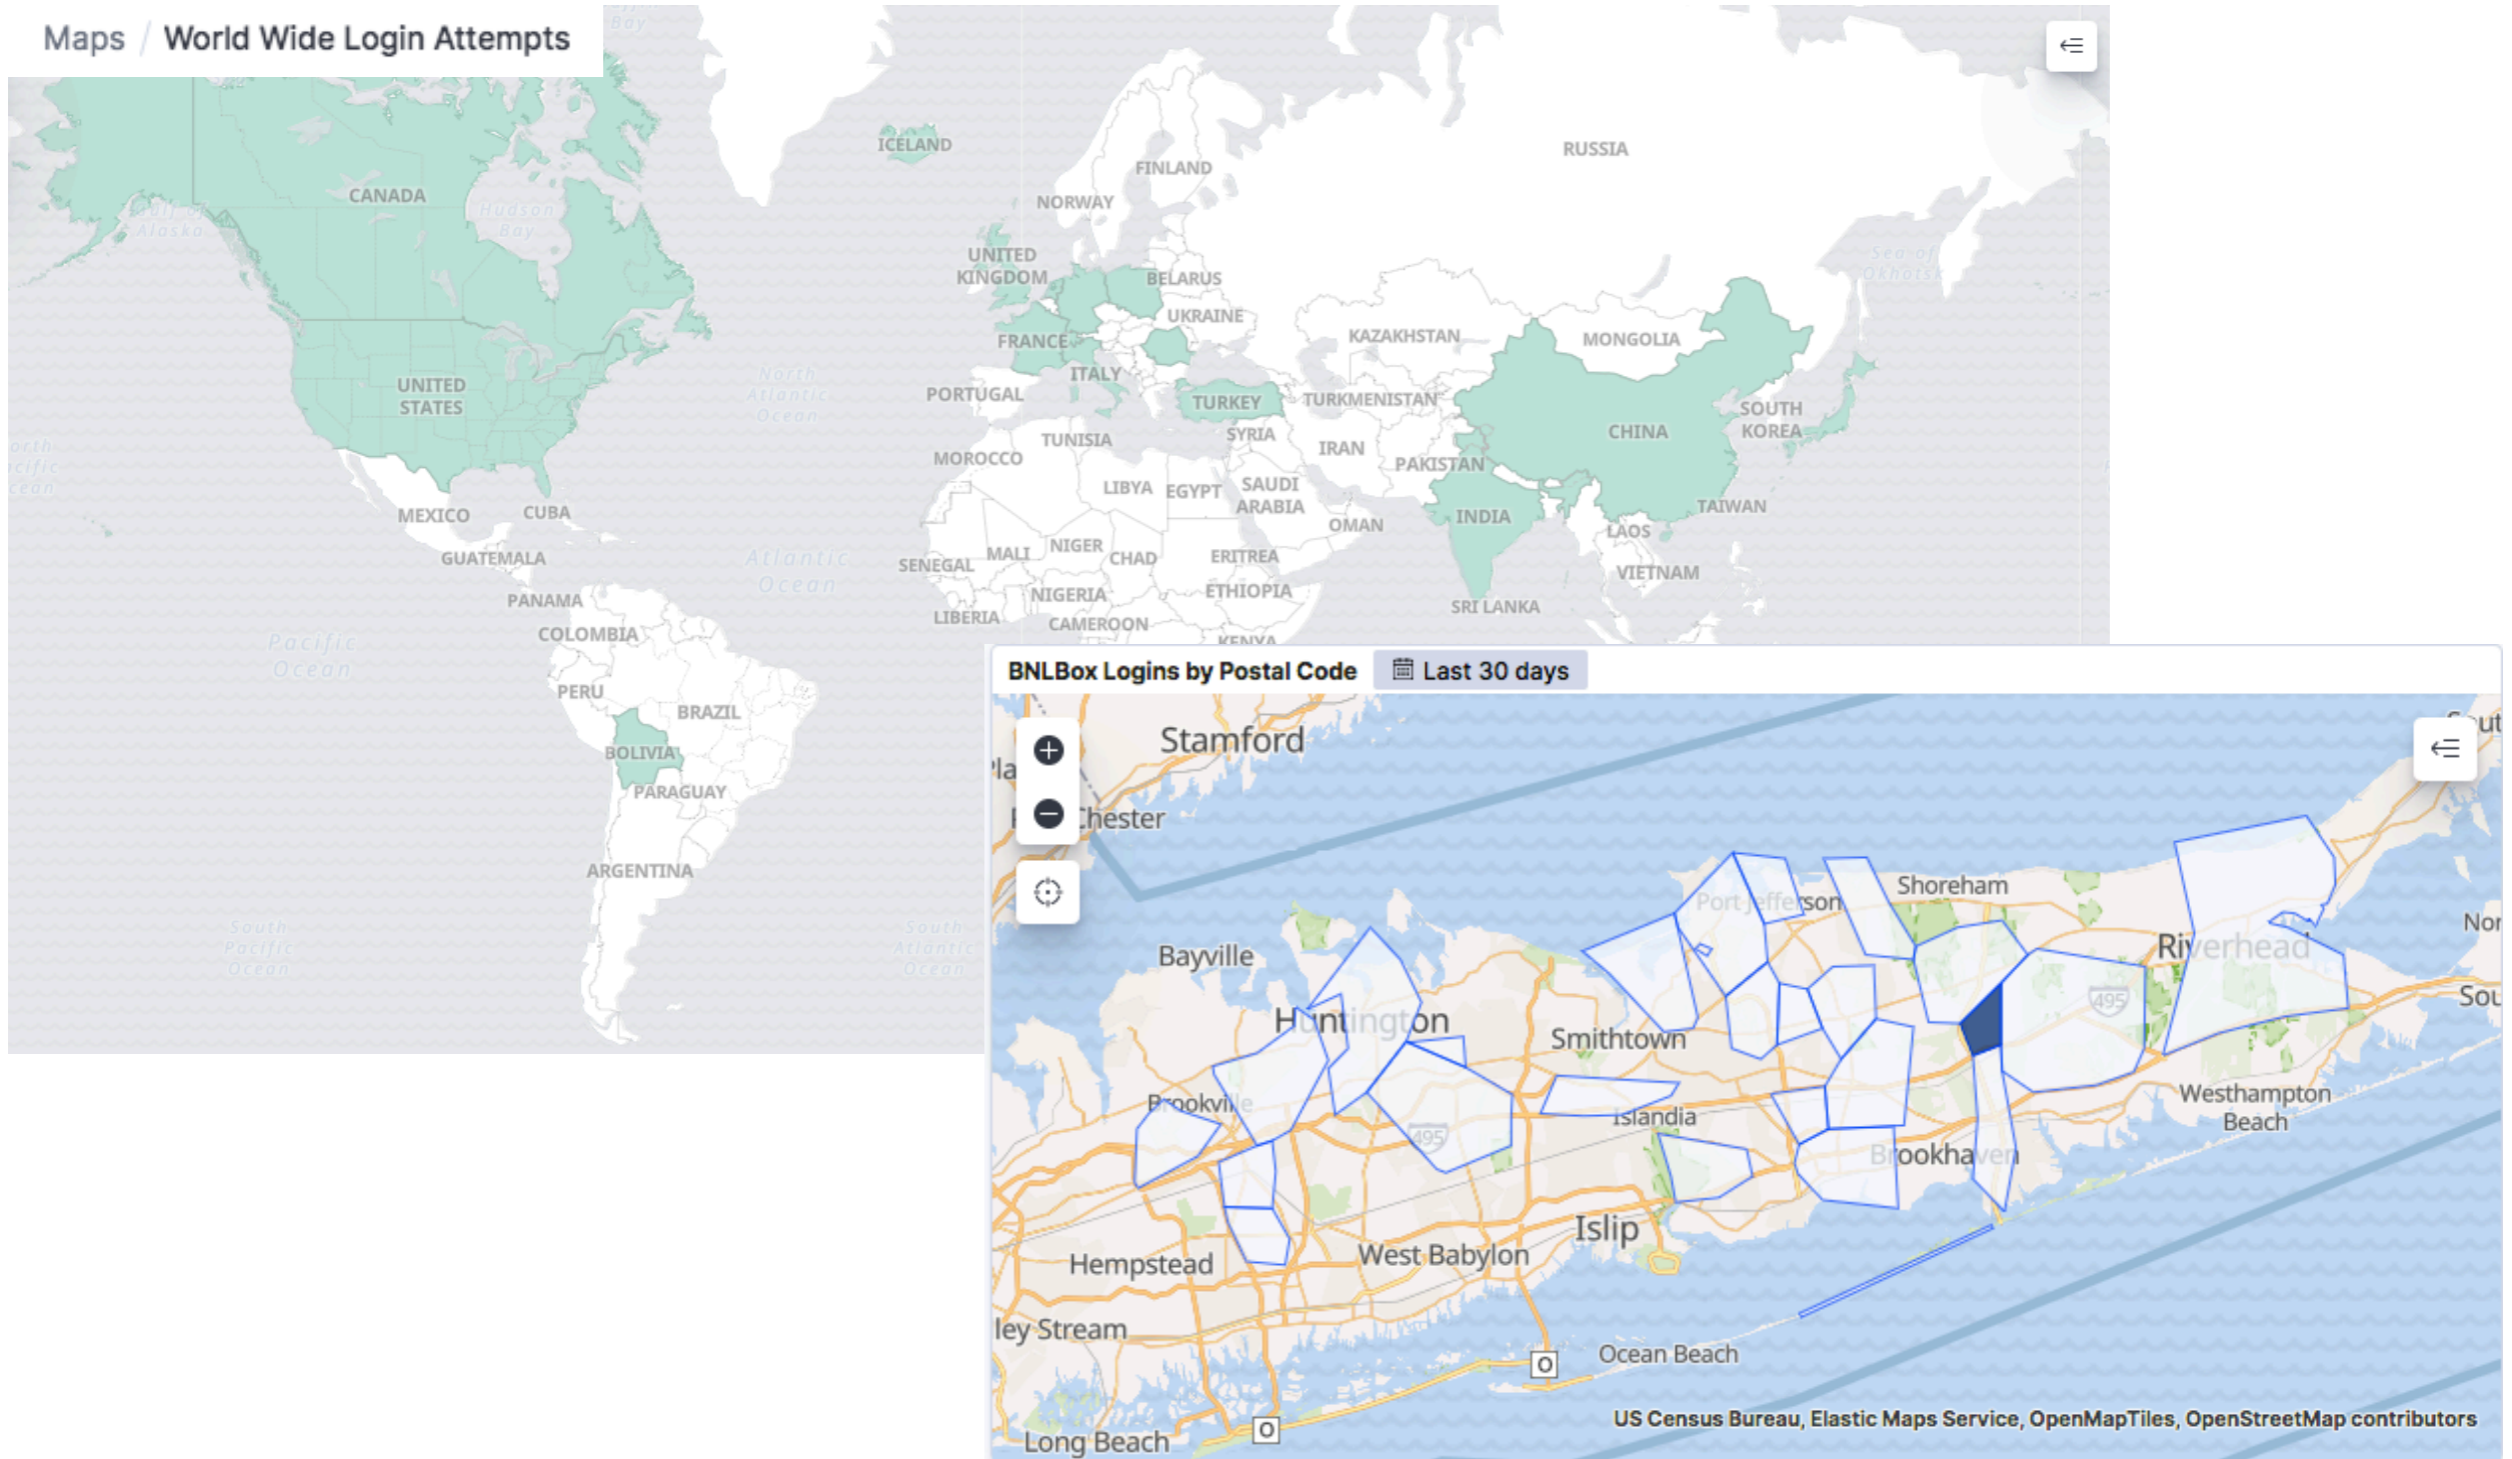
User Agents



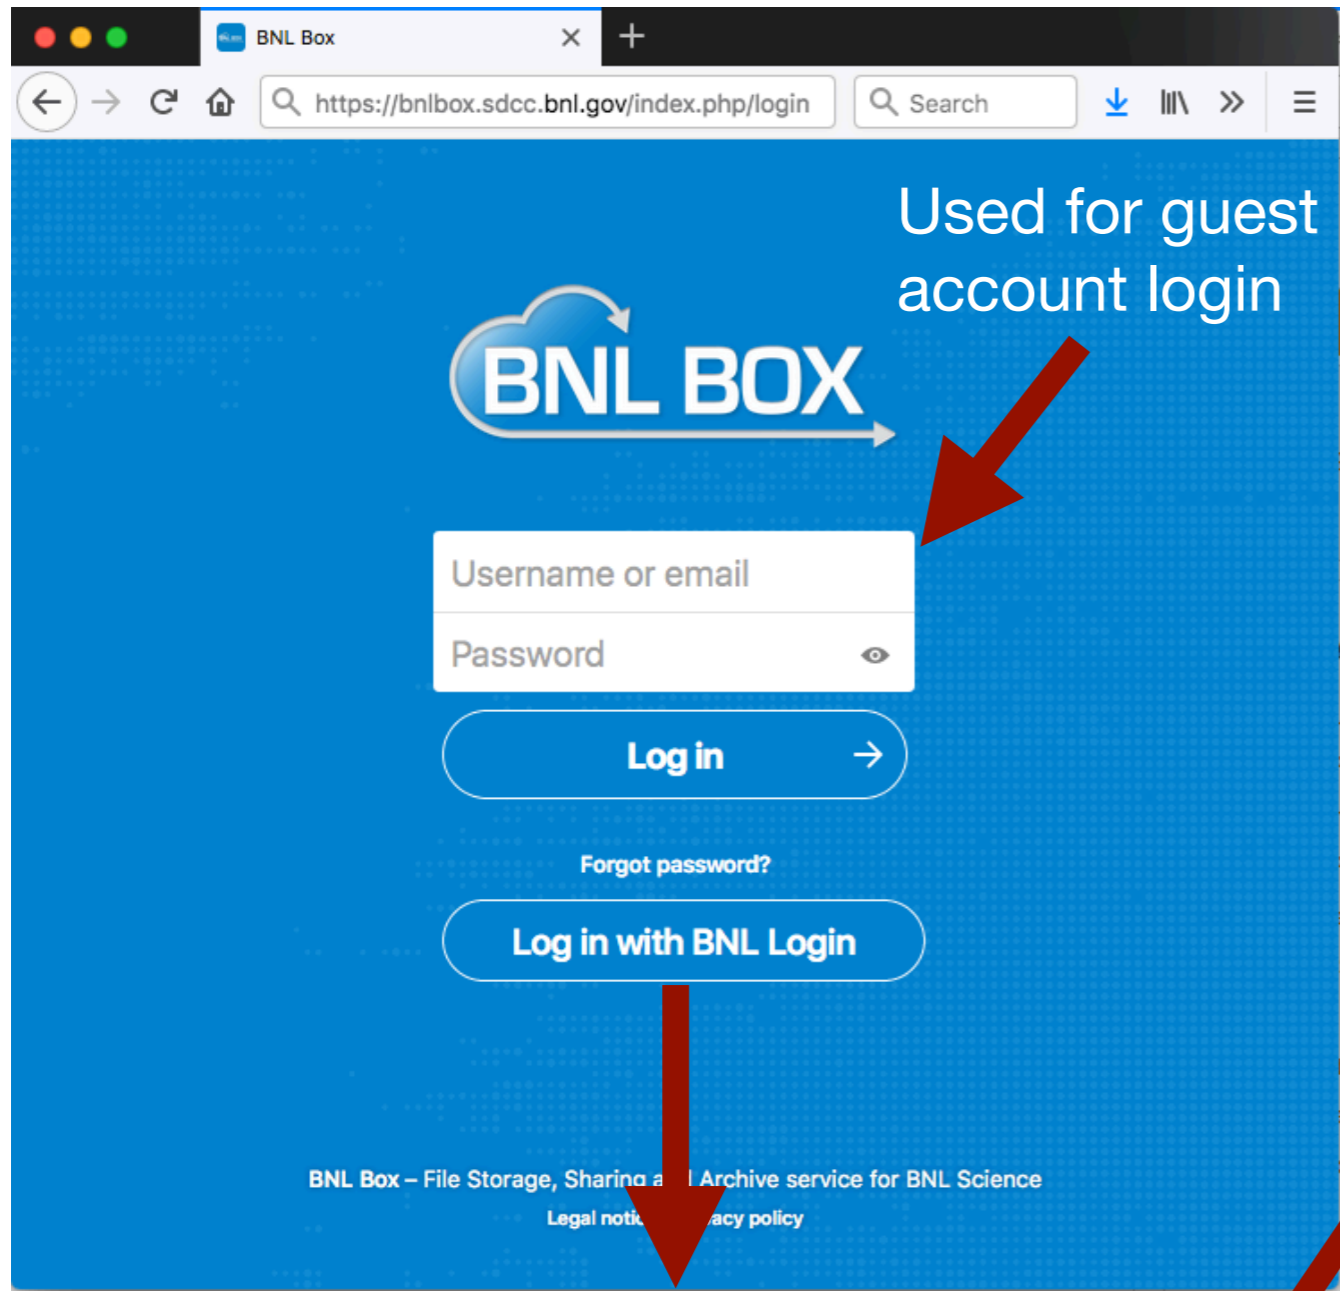
File Shares by User



GeoIP Monitoring



Features/Solutions

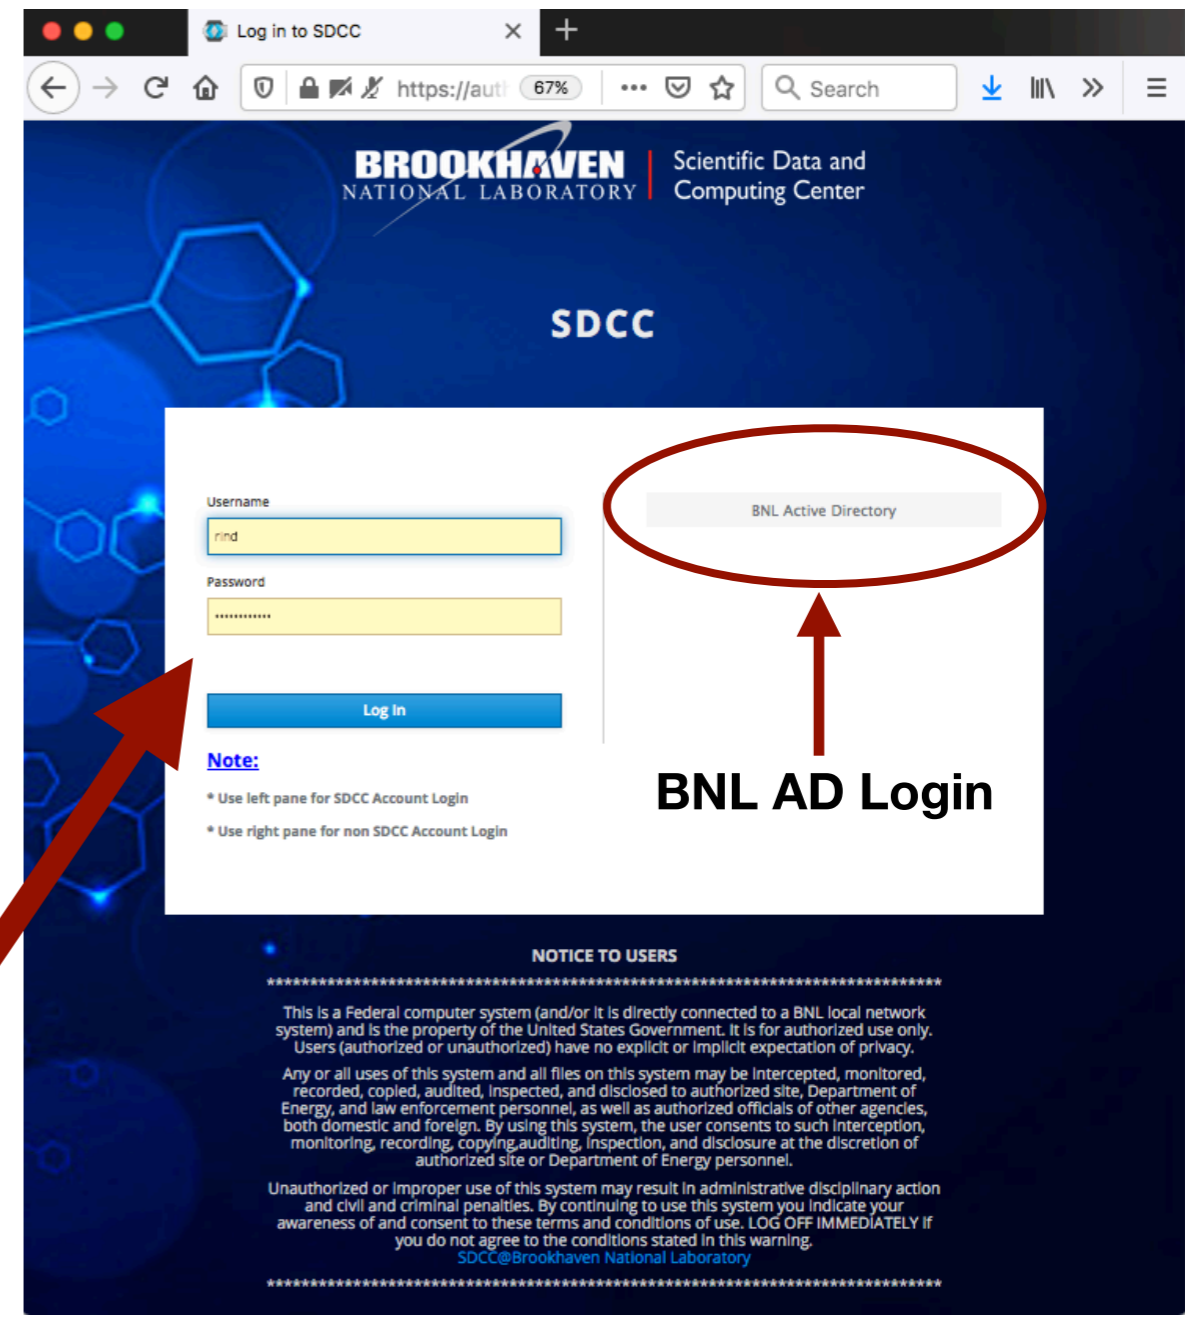


Used for guest account login

Login via Keycloak OIDC

Automated account creation upon first access
Unique UUID avoids username clashes
Duplicate prevention keyed to email address

Uses Nextcloud [SocialLogin App](#)



BNL AD Login

Option of enabling MFA and/or federated account access in future

In many cases, e.g. storing and sharing output generated by jobs running on the batch queues, users may want to transfer data to/from BNL Box in a scripted manner from the Unix command line

- Secure access provided via configurable WebDAV app
- A number of standard CL tools, such as curl, are available
- A custom ruby script, developed by H. Ito, provides a flexible, feature-rich copy tool that can handle bulk directory copies and wildcarding

For detailed documentation, see: [How To Access BNL Box From The Command Line](#)

```
[rind@rcas2501 ~]$ ./copy_bnl_box.rb -h
Usage: copy_bnl_box.rb [options] source destination
NOTE: to specify the bnlbox storage, you need to prepend 'box:' to your path seen in your bnl box service
  -h, --help                Display help message
  -s, --https_proxy ProxyHost Set https proxy
  -r, --url URL              Set none-default url
      --password password   Set User Password
  -u, --username username   Set User Name
  -f, --file filename       Configuration File
  -v, --verbose              Verbose messages
  -d, --debug                Debug option
  -p, --progress            Show progress
[rind@rcas2501 ~]$
[rind@rcas2501 ~]$ ./copy_bnl_box.rb -p box:/Photos/*.jpg .
Downloading Coast.jpg: 100% of total 819766 bytes      Done
Downloading Hummingbird.jpg: 100% of total 585219 bytes      Done
Downloading Nextcloud%20Community.jpg: 100% of total 797325 bytes      Done
Downloading Nut.jpg: 100% of total 955026 bytes      Done
Downloading Rainbow4862.jpg: 100% of total 1787072 bytes      Done
Downloading Pan_8478-81.jpg: 100% of total 11600097 bytes      Done
[rind@rcas2501 ~]$
```

Provide users with independent entry point to sharable filesystems from linux farm, non-SDCC storage, experiment online storage, etc.


External storages *i*

External storage enables you to mount external storage services and devices as secondary Nextcloud storage devices. You may also allow users to mount their own external storage services.

Folder name	External storage	Authentication	Configuration	Available for
<input type="checkbox"/> Archive	Local	None	/sdcc/lustre01/BNLBc	All users. Type to sele
<input type="checkbox"/> EIC_Data	SFTP	RSA public key	<input type="text" value="localhost"/> <input type="text" value="/gpfs02/eic/DATA"/> <input type="text" value="eickolja"/> <input type="text" value="ssh-rsa AAAAB3NzaC"/> <input type="text" value="1024"/> <input type="button" value="Generate keys"/>	Kolja Kauder ✕

NB: Files on external storage do not count against Nextcloud user's quota

- External Storage app enables mounting of shareable volumes via numerous protocols, e.g. SFTP, WebDAV, etc.
- Admin-created mounts are *read-only* and linked to Nextcloud user via PKI
- Files on external storage owned by original user, who grants account access to the Nextcloud user
- Nextcloud user may share/access external files in same manner as primary storage



- Provide long-term archival storage of large shared data sets
- Free up disk space with option of retrieval in (distant) future.
 - Archived files should be retrievable (with some latency) in transparent fashion, while not counting against user's storage quota.

- Solution: Archive Directory

- Independent directory on local Lustre filesystem mounted as External Storage
- File archiving via Lustre Copytool interface to HPSS tape storage API (configurable, driven by watermarks)
- Provisioned automatically upon BNL Box account creation and linked to user via Postgres DB trigger

- Security concerns due to ease of access from diverse entry points and sharing with unregistered users
 - Standard CyberSec port scans and log monitoring
 - Custom user agreement including appropriate file content warning (sensitive, classified, OOU, PPI, etc.)
 - Investigating AV scanning
- Formulating pricing model for power users (non-NPP)
- Exploring further integrations (fuse mounts?)
- Nextcloud Hub 20 Release

- BNL Box is a flexible, locally-hosted, secure, cloud storage solution for science, currently available to everyone with a BNL computer account
- BNL Box provides a feature-rich platform for organizing and sharing all types of data, including those generated and stored at SDCC and user facilities
- The SDCC is continuing to explore further integrations and use cases for BNL Box, driven by stakeholder needs
- Interested in other site experiences, federation, etc.

Thank You
Questions?