BNL BOX Cloud Storage Service

O. Rind, H. Ito, G. Che, T. Chou, R. Hancock, M. Karasawa,
Z. Liu, O. Novakov, T. Rao, Y. Wu, A. Zaytsev

Brookhaven National Laboratory

BROOKHAVEN NATIONAL LABORATORY | Scientific Data and Computing Center

Overview

"Dropbox" style file syncing/sharing service integrated into the BNL SDCC to provide an easy-to-use, unified storage service for all BNL users

- Locally-hosted cloud storage, providing easy access to user data via browser, desktop & mobile clients, including automated synchronization
- Flexible file-sharing methodologies for sharing data with external collaborators
- Multiple access points from local user accounts, including from batch system, analysis portal, etc.
- Tape archiving capability (via Lustre & HPSS) for large data sets
- Open to all users with SDCC or BNL AD accounts (via Keycloak based SSO) with straightforward capabilities for Federated ID
- An integrated effort by many SDCC staff supporting back-end disk and tape storage, web and db services, AAI front-end, user interface, etc.

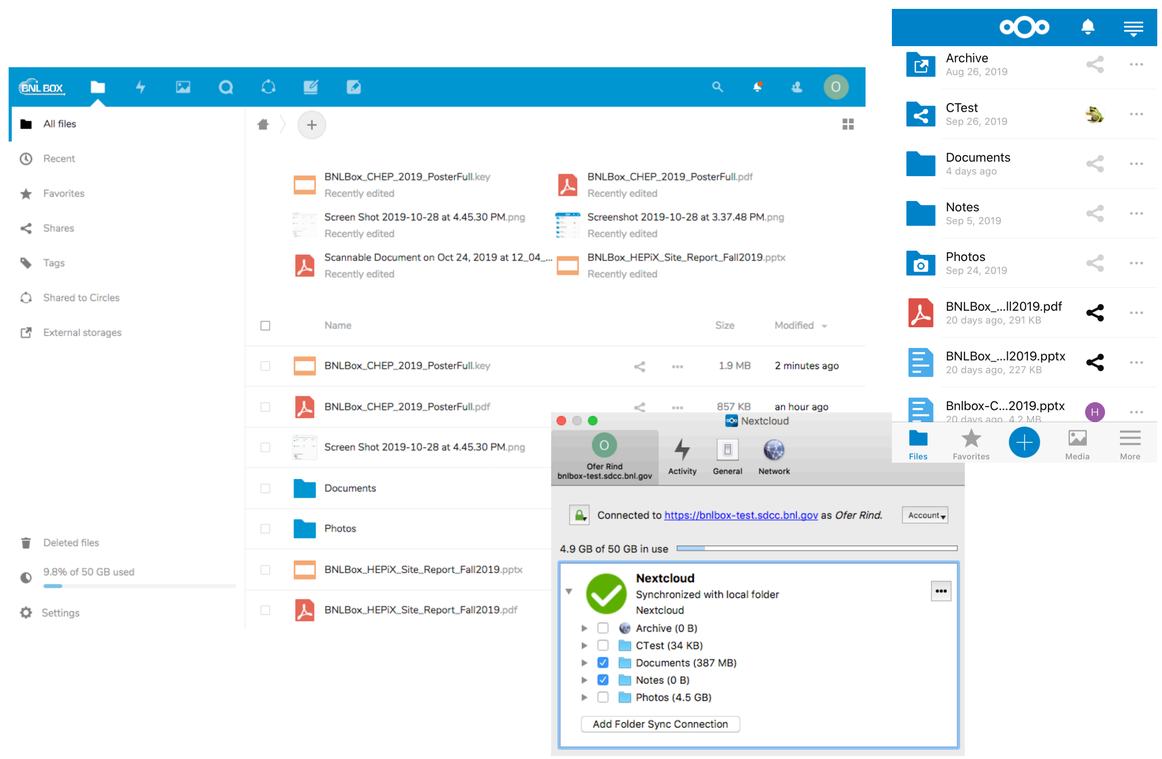

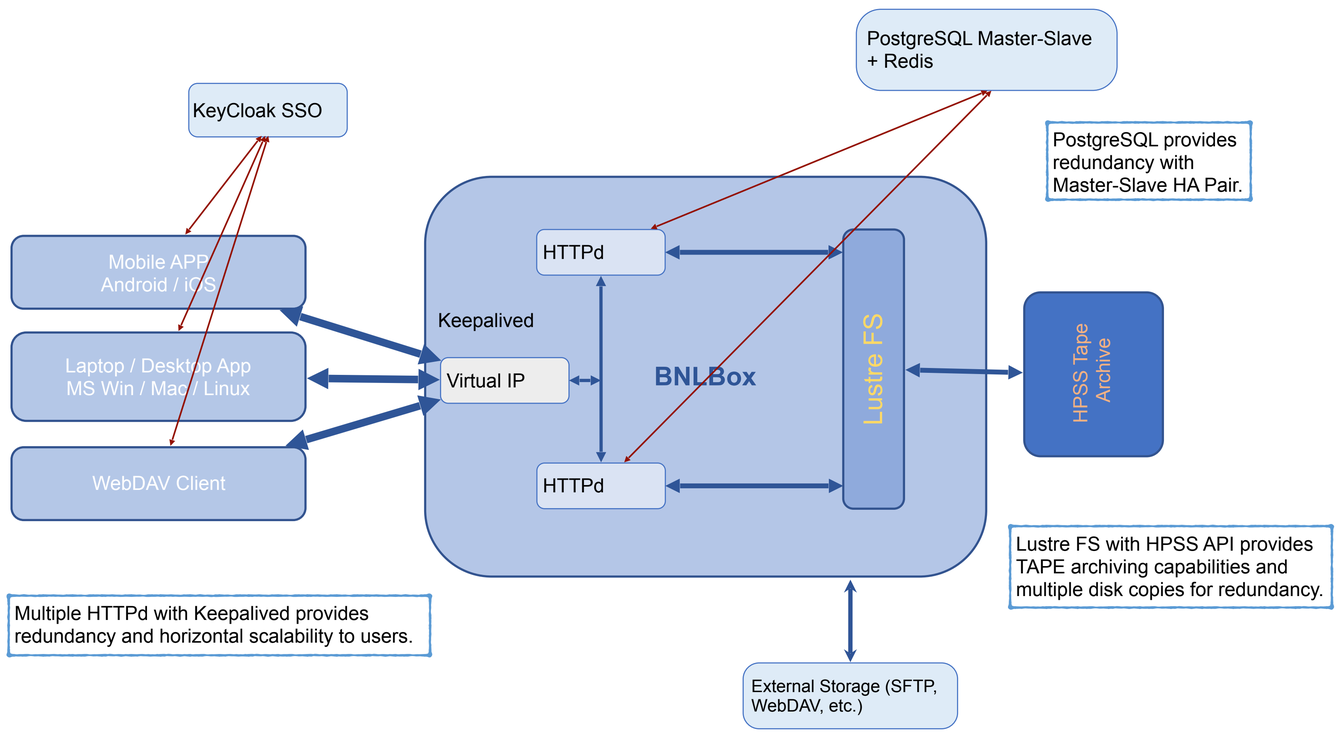

Components

- Former Owncloud-based BNL Box service has been completely redesigned using the latest open source Nextcloud (v16)
- All-new production level server and storage hardware
  - Load-balanced, HA-pair server architecture
    - Load balancing via round-robin DNS + keepalived
    - Shared data storage and configuration directory
    - Shared PSQL and Redis database backend (with standby)
  - Resilient, high-performance Lustre file storage with TSM backup
  - HPSS tape archive

Features: AAI

- BNL Box accounts linked to SDCC/BNL auth infrastructure via Keycloak and Social Login app Custom OIDC configuration
- Accounts created according to unique Keycloak uuid - ensures no UID conflicts from independent user databases
- Single email enforced to prevent multiple accounts
- 2FA enabled
- Easy integration of federated accounts, e.g. CILogon
- No local accounts other than admin or guest

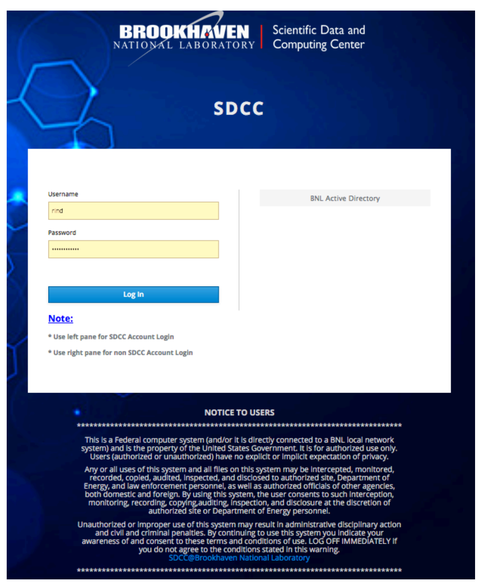

Features: Archive

Users have requested long-term archival storage of large shared data sets - need mechanism to 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.

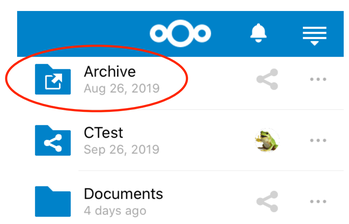

- Independent Lustre directory mounted as External Storage
- Provisioned automatically upon BNL Box account creation and linked to user via Postgres trigger
- File archiving via Lustre Copytool interface to HPSS tape storage API

Features: External Storage

Provide users with independent entry point to sharable filesystem from analysis farm, batch system, non-SDCC storage, experiment online storage, etc.

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

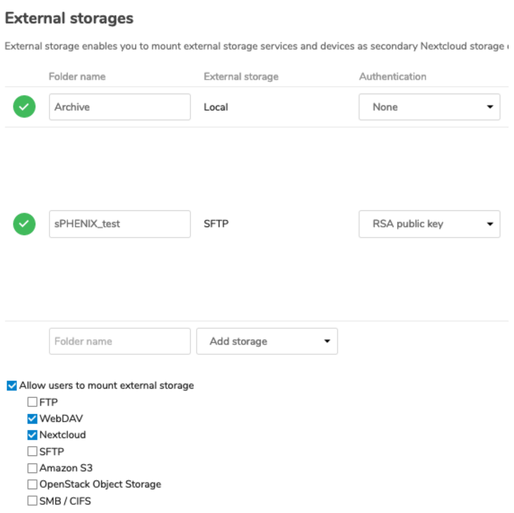

Issues & Concerns

- Transition from former Owncloud-based systems requires users to migrate their data
  - Temporarily allow user mounts of Owncloud storage via WebDAV to copy folders during migration period
- Shared storage via anonymous link or guest account raises security concerns
  - Enforce read-only link sharing
  - Log analysis/virus scanning
  - Visible user warnings/disclaimers
- Federated user access requires an updated Computer Use Agreement
- Handling of Nextcloud auth tokens