BNLBox Status Report

Ofer Rind, Hironori Ito

Brookhaven National Laboratory

BROOKHAVEN NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN SCIENCE ASSOCIATES

BNL and Cloud Storage Requirements

- **BNL**: **Brookhaven National Laboratory**
  - Multi-purpose US Department of Energy National Laboratory located in NY
  - Conducts a wide range of scientific activity: medical and biological studies, chemistry and nano-material studies, powerful light source, nuclear and particle physics.
  - Hosts Relativistic Heavy-ion Collider, **RHIC**.
  - Will host Electron-Ion Collider, **EIC**.
  - Hosts National Synchrotron Light Source II, **NSLS-II**
  - Hosts National Center for Functional Nano-materials, **CFN**

- **Cloud Storage Requirements**
  - Allow not only nuclear and particle physics BNL users to have cloud storage, but all BNL scientific and non-scientific users
    - KeyCloak with OpenID
  - Allow BNL users to share with non-BNL users
    - Guest features
  - Allow users to create and manage their own groups
    - Circle features
  - Provide archival storage
    - Lustre with external storage

Status and Key features

- **BNLBox** became a production service of the SDCC in Nov 2019.
  - It replaced the test setup running for the last 2 years.
- It is based on Nextcloud. It is currently v17.
- The database is using PostgreSQL
  - Master and slave for redundancy setup.
  - Fast NVMe disk for performance.
  - PostgreSQL will allow us to partition the tables for further performance tuning when necessary.
- The file system is Lustre file system.
  - Multiple copies for redundancy.
  - HPSS API for archival storage.
  - 1PB of usable storage
  - Backed up by TSM.
- KeyCloak is being used for login authentication.
  - OpenID connection with KeyCloak via Social Login.
  - Integrates SDCC LDAP accounts as well as BNL Active Directory.
  - Allows any BNL users to create accounts.
- Multiple front-end Apache servers
  - KeepAlived provides the redundancy and load-balancing.
  - 40 Gbps WAN-to-storage currently.

Components

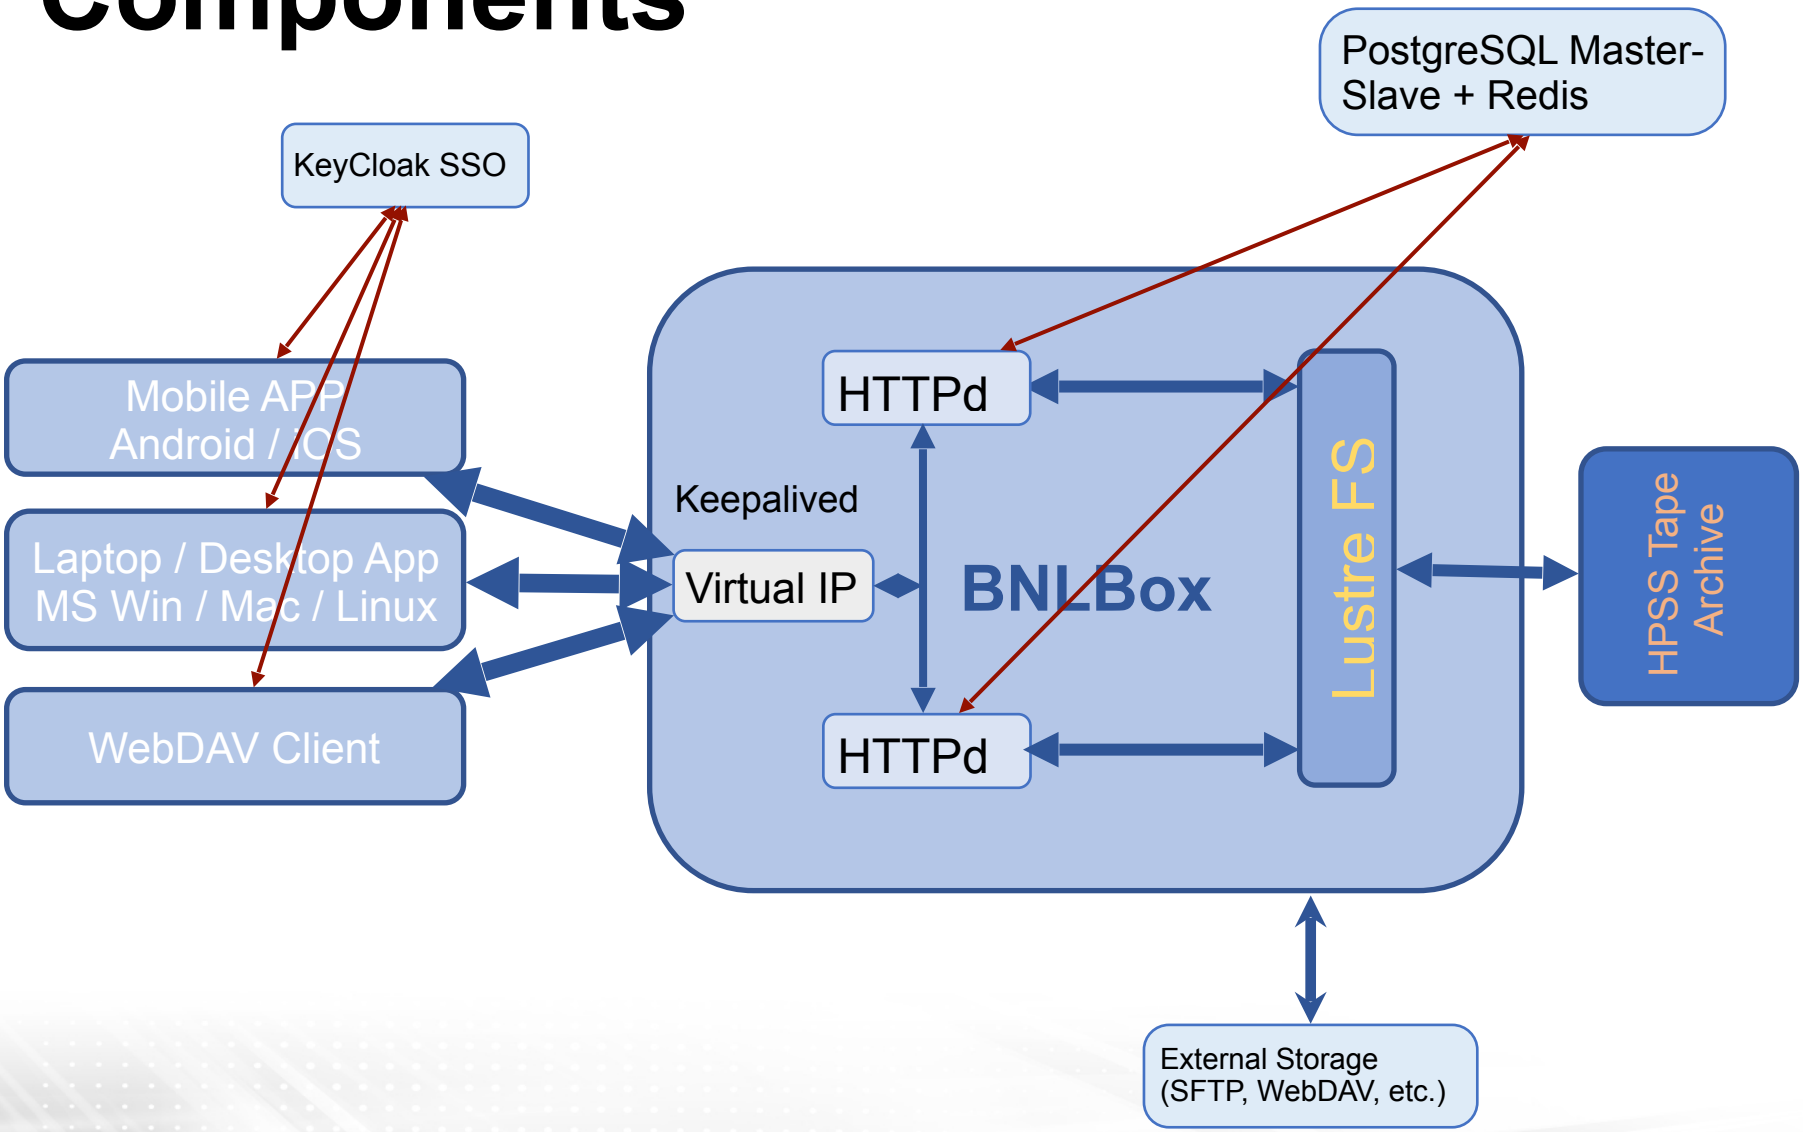
PostgreSQL Master-Slave + Redis
KeyCloak SSO
Mobile APP Android / iOS
Laptop / Desktop App MS Win / Mac / Linux
WebDAV Client
Keepalived
Virtual IP
HTTPd
HTTPd
BNLBox
Lustre FS
HPSS Tape Archive
External Storage (SFTP, WebDAV, etc.)

Current use

- Since the production release in November, a few users are being added every day, resulting in ~100 users right now.
- The traditional nuclear and particle physics users are interested in the sharing of their data and documents amongst themselves, while users of the light source and nano-center are interested in the export of their acquired data to remote users.
- Integration with BNL Jupyter is being considered.
- Sharing of large remote storage (e.g. GPFS) via external storage feature.

Missing features and concern

- WebDAV with “social login”
  - Some users with large data (~10TBs) prefer “cp” over “sync,” but with “Social Login” auth, simple WebDAV commands stopped working because the client must understand the special social login flow. The same can be said of WebDAV mount to the remote client.
- Scalability
  - In contrast to the commercial cloud services, the number of users of BNLBox will be small; a few thousand at most. But, they will likely use more space than typical users and some will use very heavily.
  - LustreFS with a lot of small files.
  - PostgreSQL
- Security
  - Users are expected/warned not to store any personally identifiable data (PII) or any other documents requiring special care and security
  - Users should not store any illegal data.