

HPSS Connections to ESGF: BASEJumper

Sam Fries¹, Sasha Ames¹, Alex Sim², Dean Williams¹

¹AIMS Group, Lawrence Livermore National Laboratory
²SDM Group, Lawrence Berkeley National Laboratory

Abstract

Raw climate model output can range in size from hundreds of terabytes to petabytes, which necessitates tape storage rather than hard disks. Retrieving the data from tape can be difficult, time consuming, error prone, and when the it is finally retrieved, it may not even be the correct dataset. To address this, we have created a bridge between HPSS and ESGF, allowing datasets stored on tape to be discovered and accessed using the faceted search tools provided by ESGF. LBNL's Berkeley Archival Storage Encapsulation (BASE) library provides a simple python API for retrieving metadata as well as actual data from HPSS; using this library, we have created BASEJumper, an application that consists of two components: a web frontend that interfaces with ESGF's services and manages HPSS transfer metadata, and a daemon that manages retrieval and storage of the data. We have also created a new publishing script that safely exposes paths from HPSS for publishing using the frontend's RESTful API along with the BASE library to extract necessary metadata for ESGF.

Web Frontend

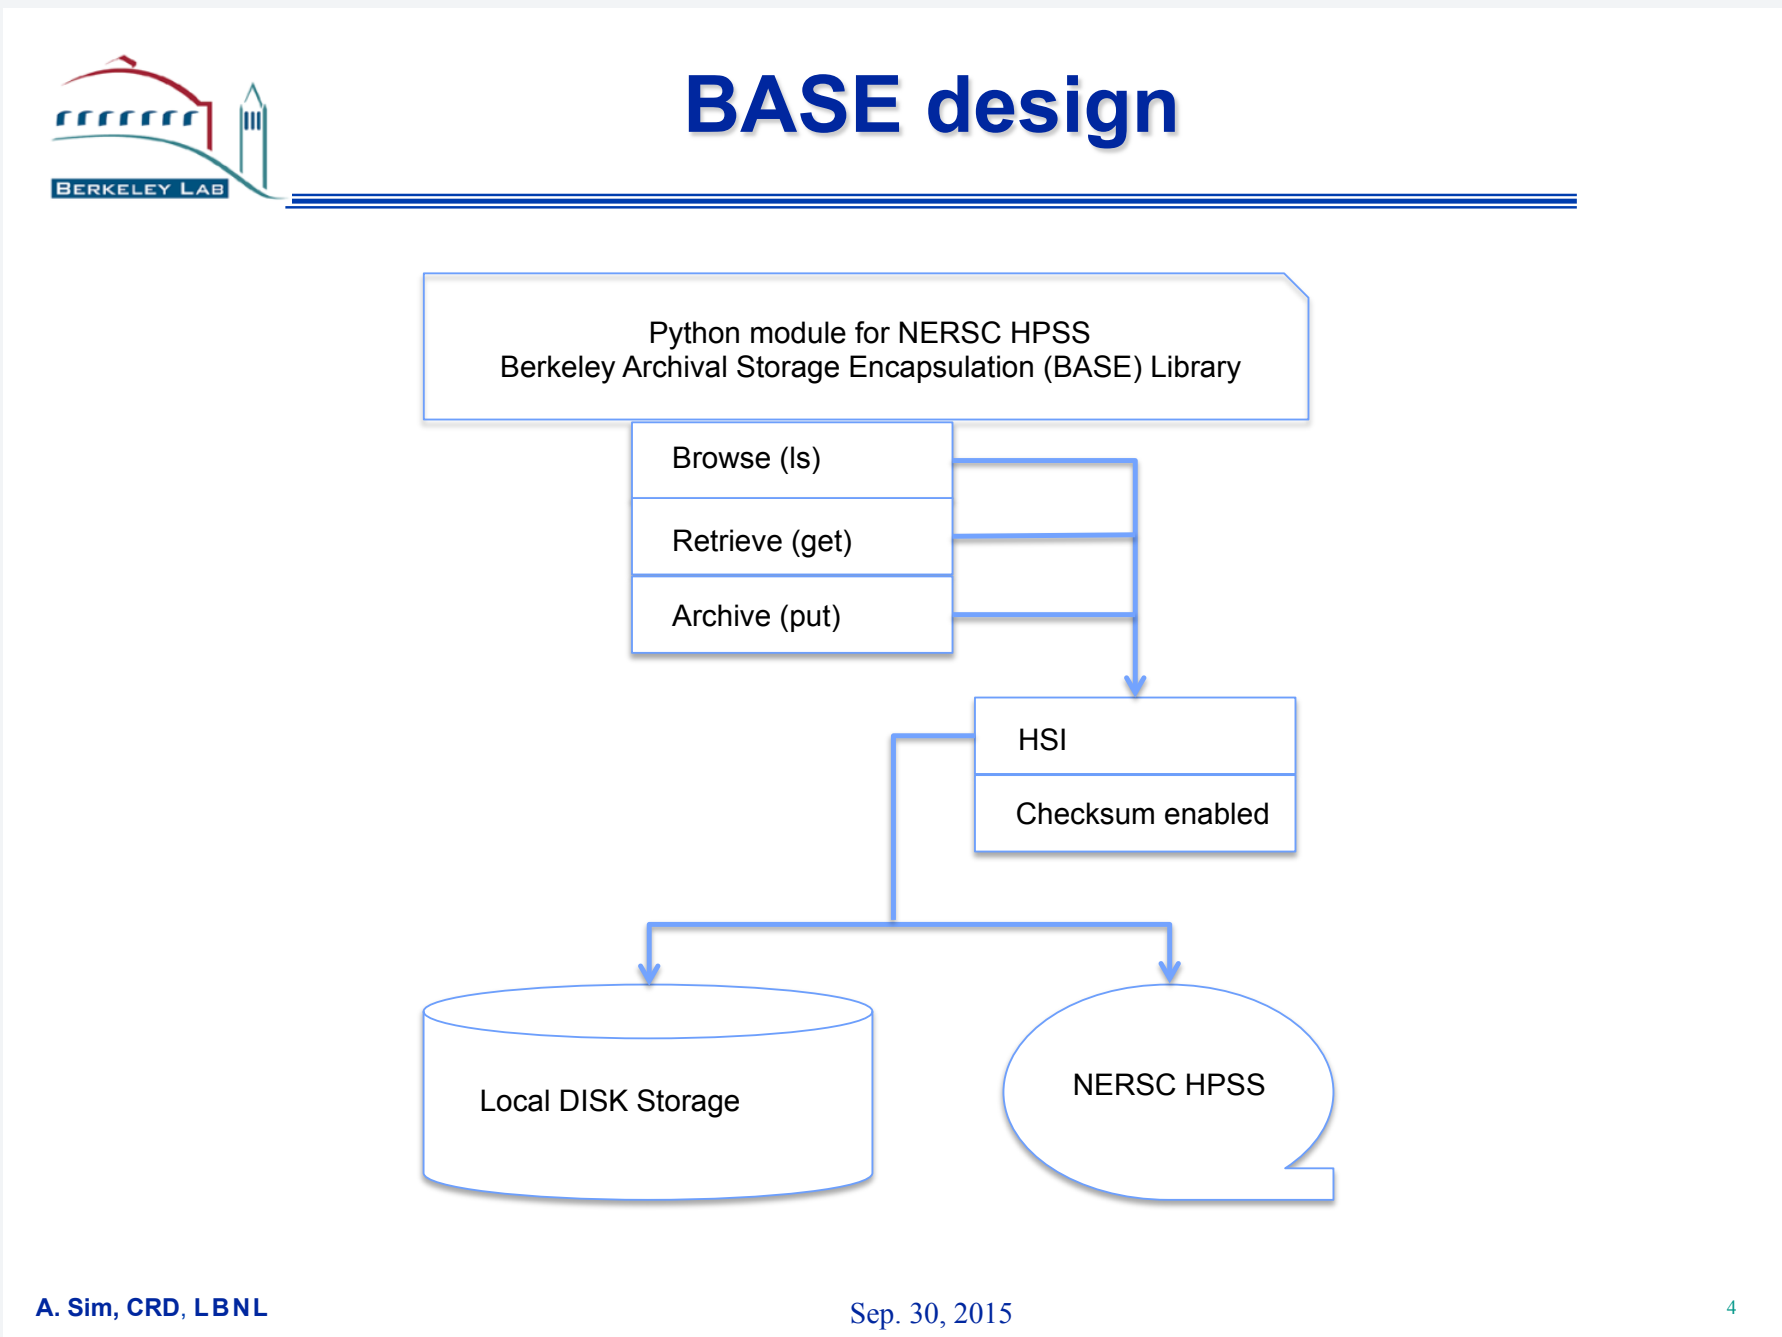
The web application is integrated with our CoG install to provide a user-friendly view into the transfer status. The publisher enters the BASEJumper files into the search catalog using the "BASEJ" moniker; if the user has access to the file (using standard ESGF group permissions), clicking on the link will queue the transfer for the daemon. The user can view their progress on their profile page in CoG. Once the transfer is complete, anyone waiting on the transfer will receive an email notification with the download link, which will also appear in their profile.

Daemon Backend

The daemon manages data retrieval from HPSS using BASE. It maintains a local cache of data transferred off HPSS, and evicts files as space is needed for new transfers. As data is received from HPSS, it updates the progress of the transfer using a RESTful API in the web frontend. Once the data is extracted from the tape archive, it is transferred to the web frontend's data cache, and a notification is sent to users that the data is available for download. Since the daemon communicates with the frontend via a RESTful API, it can be run within a firewalled enclave, to retrieve data without requiring token-based authentication. The mechanism for transferring the data from the daemon to the frontend is defined in the configurations for the daemon, so the complexity can range from a simple cp to a DTN-staged GridFTP transfer.

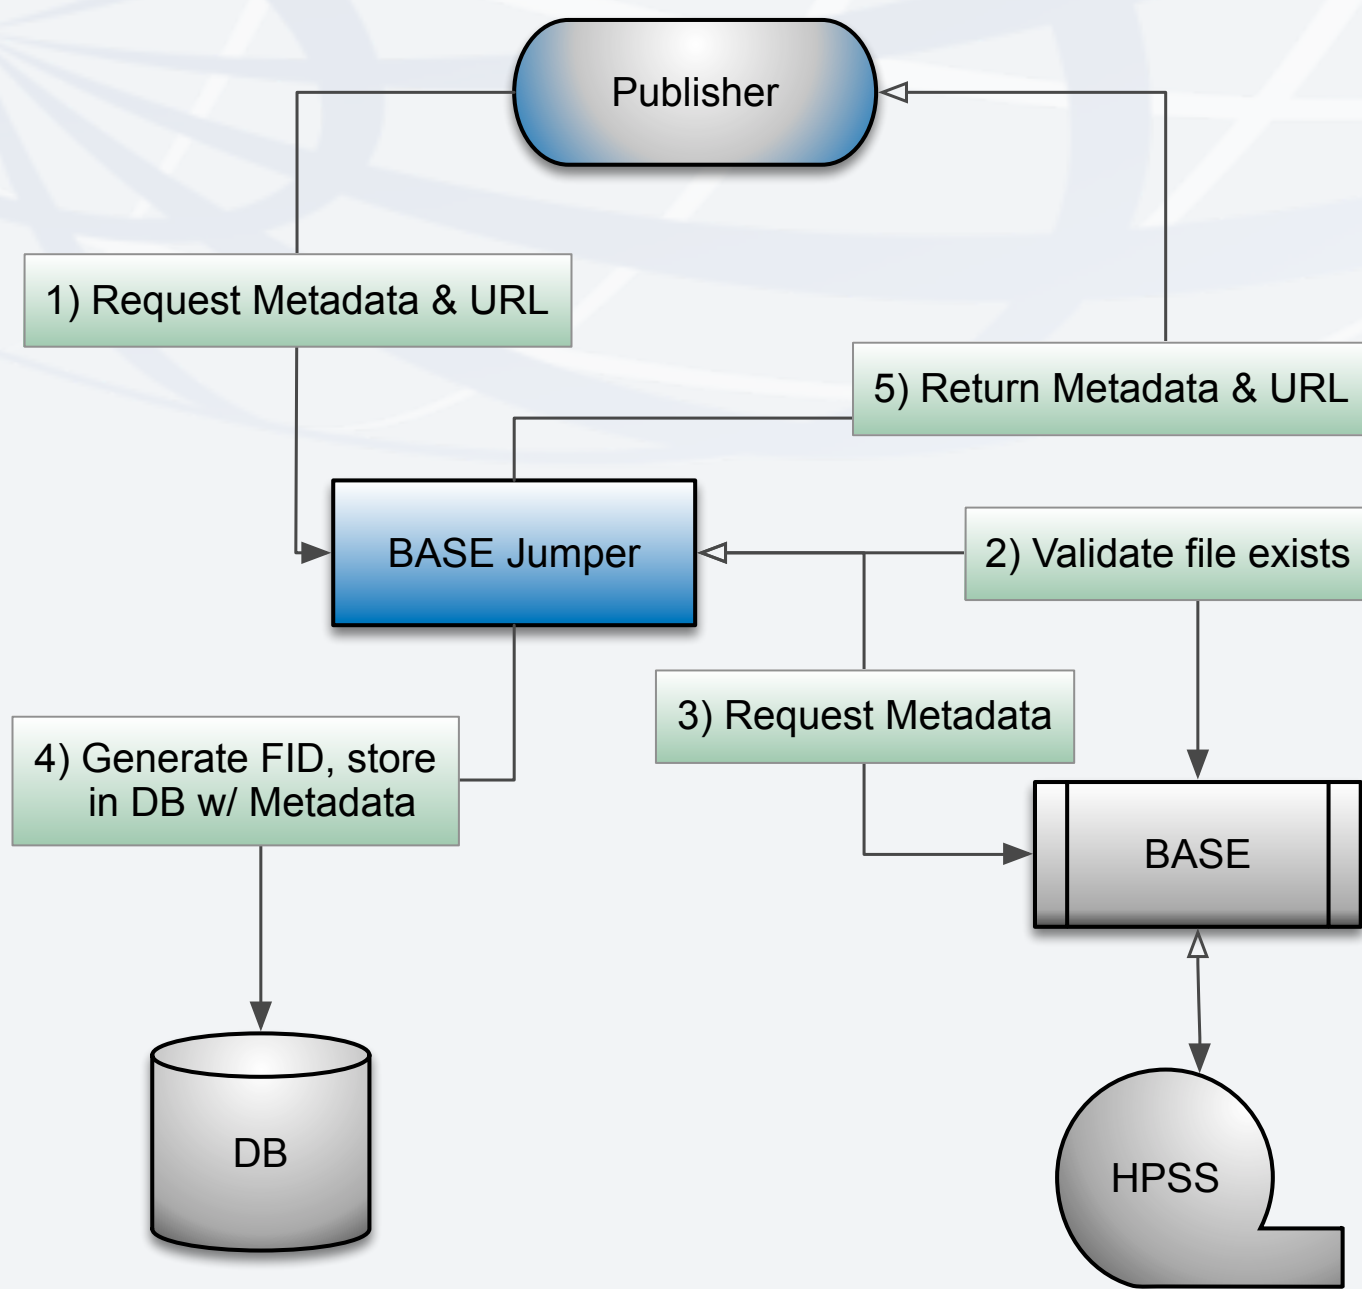
BASE

The Berkeley Archival Storage Encapsulation Library comes from the experience of building, maintaining, and operating the BeStMan and HRM systems from 1998-2016. It is a small python module that provides three main functions (browsing, retrieving, and archiving). Internally, it leverages HSI for performing these tasks, creating a greatly simplified design.



Publishing Workflow

The publisher utilizes scripts packaged with BASEJumper to retrieve metadata about files. It exposes the paths using the BASEJumper RESTful API, and publishes the URLs through the TDS catalog and SOLR. These URLs will then appear in the search results shown in CoG, from which users will gain access to each file or dataset stored in HPSS.



Future Work

There are a number of major improvements that can be made to BASEJumper to facilitate the ease of use for retrieving large volumes of data from HPSS archives.

- Publish individual files from TAR bundles
- Expose files for download via Globus/GridFTP rather than HTTP
- Add support for retrieving multiple files in one transfer
- Improved cache logic for staged files