EOSDIS
NASA'S EARTH OBSERVING SYSTEM
DATA AND INFORMATION SYSTEM

Trade Study: Storing NASA HDF5/netCDF-4 Data in the Amazon Cloud and Retrieving Data via Hyrax Server / THREDDS Data Server

Ted Habermann, Aleksandar Jelenak, Joe Lee, Kent Yang,
The HDF Group
James Gallagher, Nathan Potter
OPeNDAP, Inc.

This work was supported by NASA/GSFC under
Raytheon Co. contract number NNG15HZ39C

Goal

- Study one or more integrated solutions for storing and retrieving NASA HDF5 and netCDF4 data using Amazon Web Services (AWS) Simple Storage Service (S3) and the Hyrax server.
- Explore strategies for granulizing and aggregating data that optimize both performance and cost for data storage and retrieval.
- Develop a cloud cost model for the preferred data storage solution that accounts for different granulation and aggregation schemes as well as cost and performance trades.

Methodology

- Three different architectures to study.
- Three different NASA data collections uploaded to S3.
- Index file content and dataset byte storage information.
- Seven use cases.

Architecture #1: Baseline Hyrax Data Access

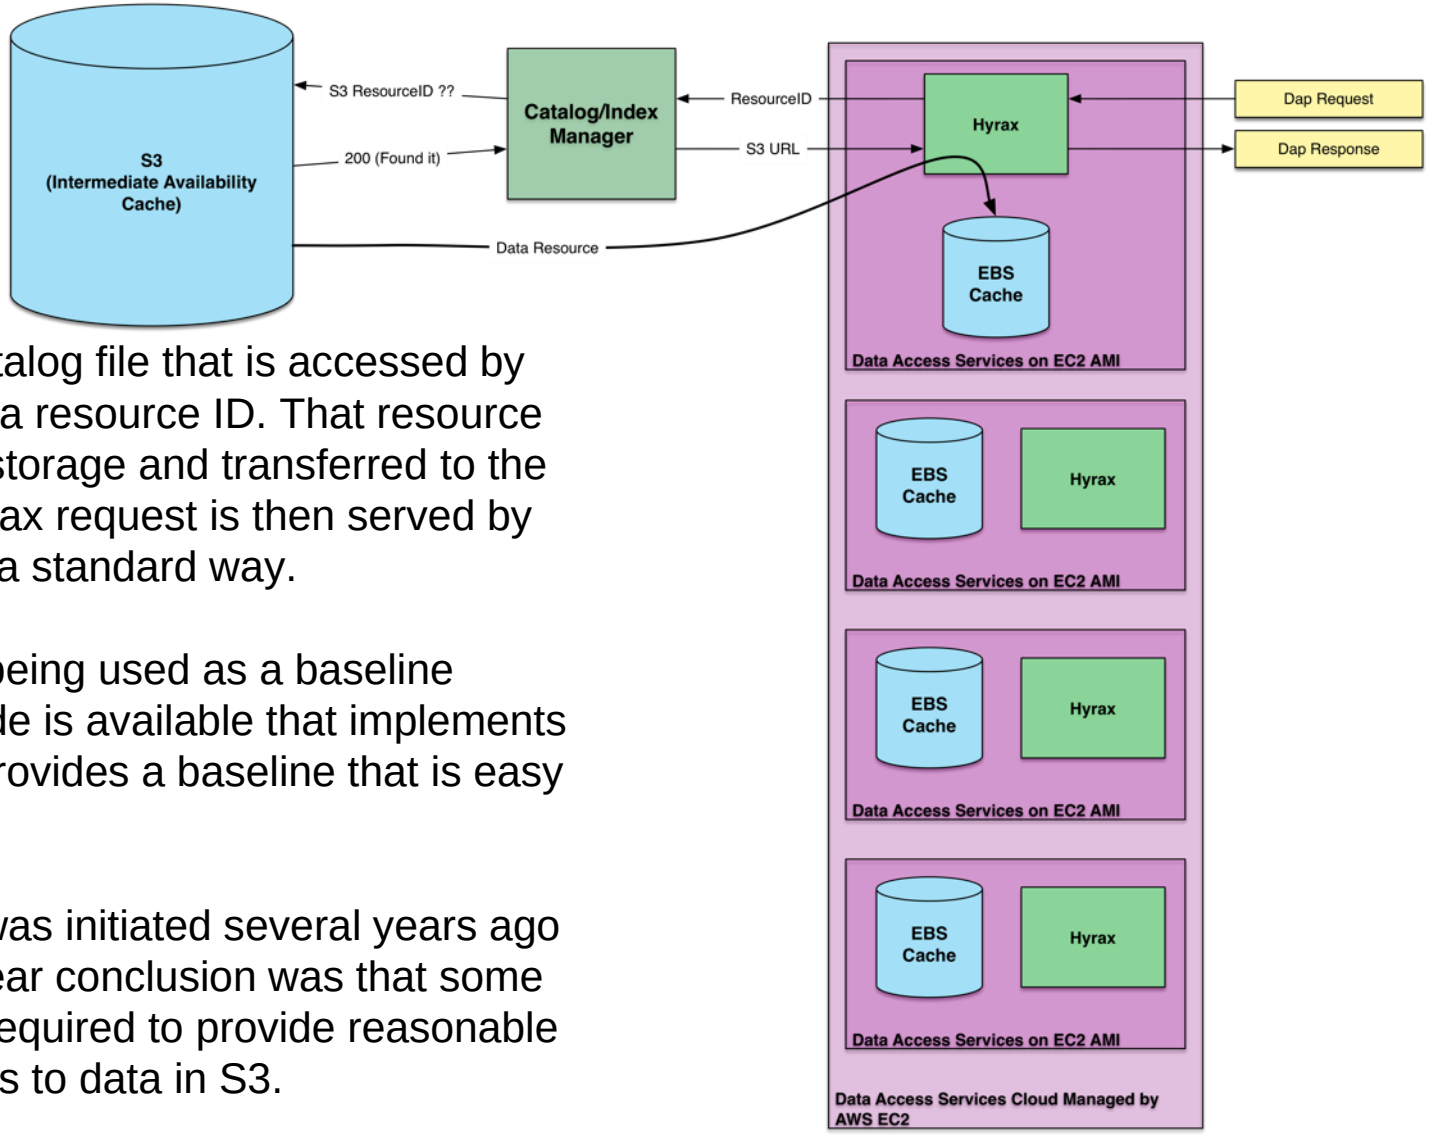

Data in S3 with a catalog file that is accessed by the Hyrax server for a resource ID. That resource is located in the S3 storage and transferred to the EBS cache. The Hyrax request is then served by the EC2 machine in a standard way.

This architecture is being used as a baseline because running code is available that implements it and, therefore, it provides a baseline that is easy to get going quickly.

Note that this work was initiated several years ago with NOAA and a clear conclusion was that some sort of catalog was required to provide reasonable discovery and access to data in S3.

EOSDIS

Architecture #2: Files With Range-Gets

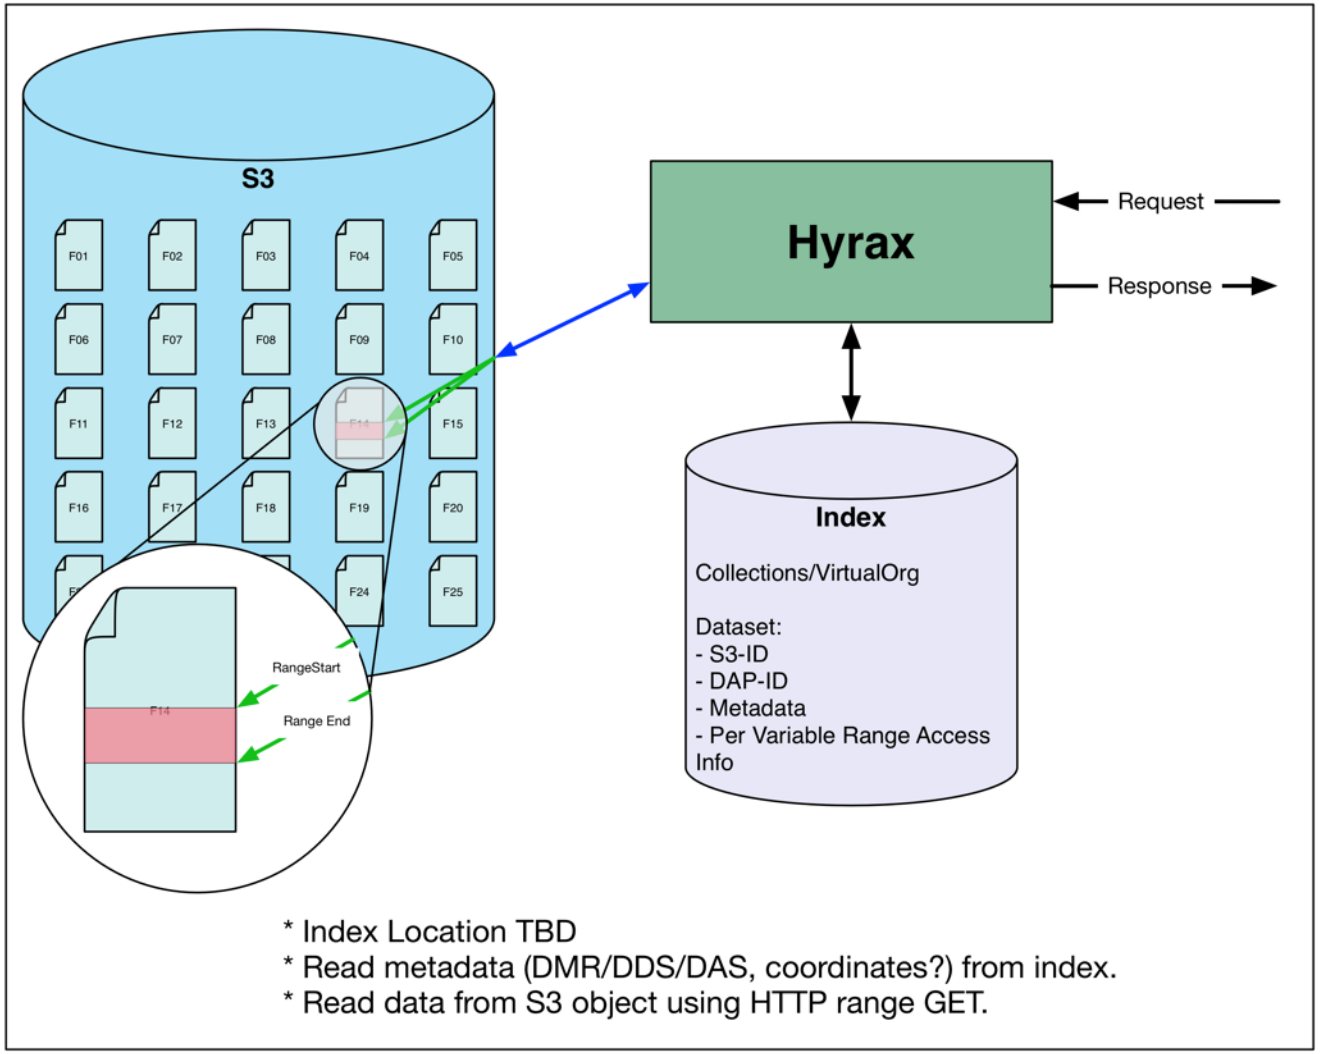

Architecture #3: HDF5 Datasets as Objects

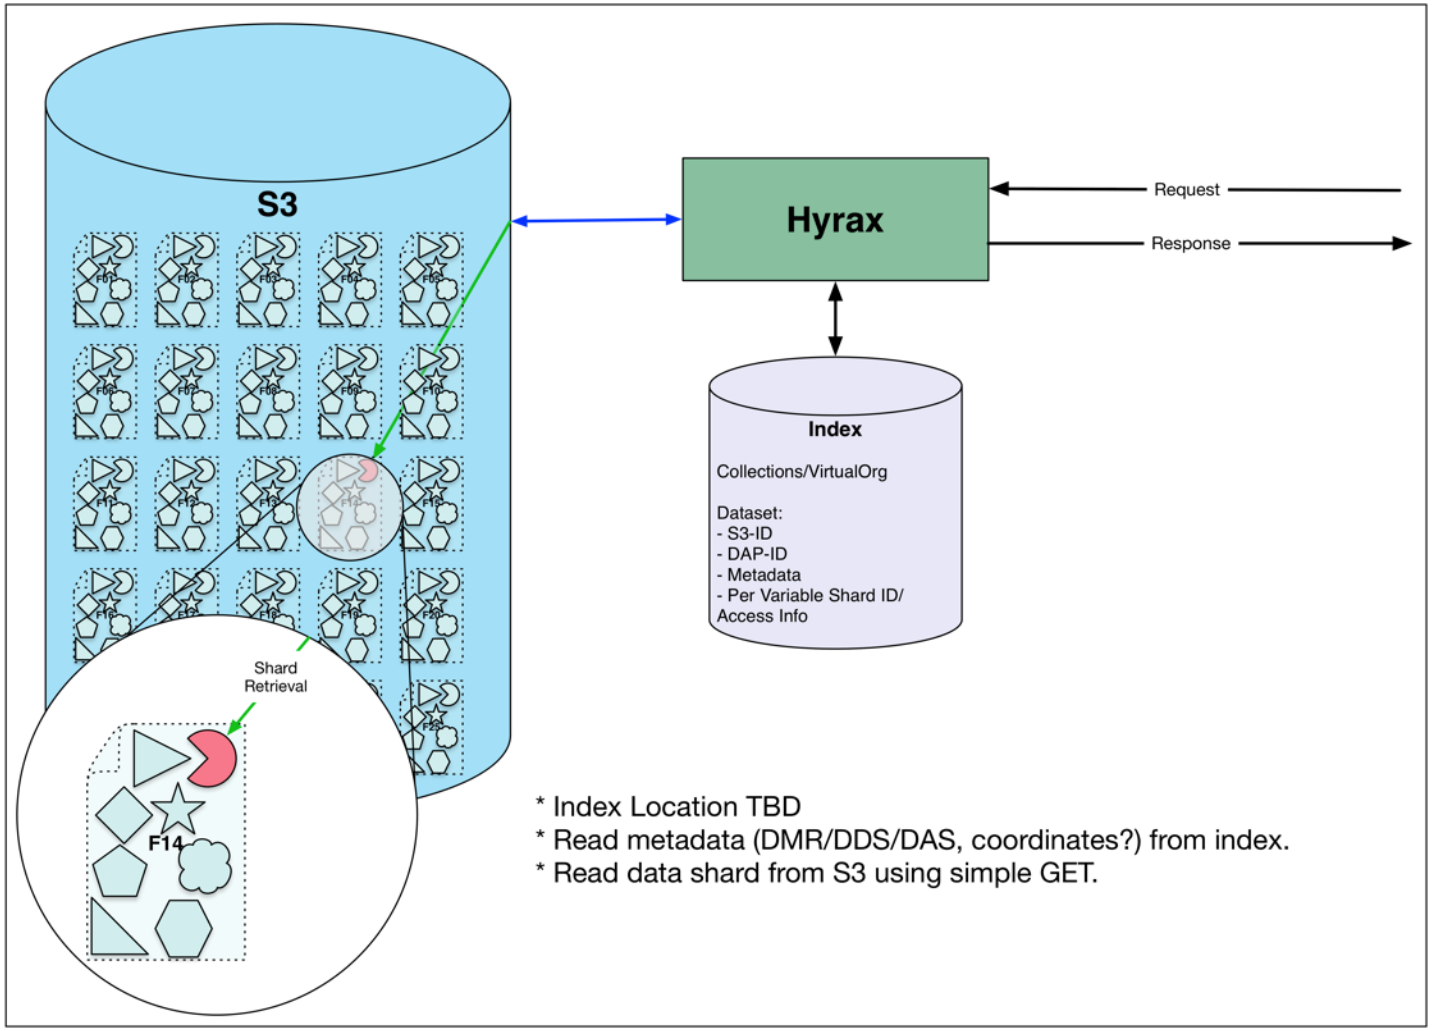

Data Collections in AWS S3

- NASA AIRSSTD3.006 product for year 2015 (365 files; 122GB)
- NASA MERRA2 product (7,466 files; 231GB)
- Sample HDF5 files (96 files; 125MB)

Index Files

- A catalog of file content and dataset byte storage information.
- One for each file in the data collections.
- Hyrax's Dataset Metadata Response (DMR) XML used as the basis.
- HDF4 File Map XML used for dataset storage information.

Index File Content

```
<?xml version='1.0' encoding='UTF-8'?>
<Dataset xmlns="http://xml.opendap.org/ns/DAP/4.0#"
        xmlns:h4="http://www.hdfgroup.org/HDF4/XML/schema/HDF4map/1.0.1"
        dapVersion="4.0" dmrVersion="1.0">
  <Dimension name="Latitude" size="180"/>
  <Dimension name="Longitude" size="360"/>
  <Float32 name="ClrOLR_A">
    <Dim name="/Latitude"/>
    <Dim name="/Longitude"/>
    <h4:chunks deflate_level="2" compressionType="shuffle deflate">
      <h4:chunkDimensionSizes>180 360</h4:chunkDimensionSizes>
      <h4:byteStream nBytes="72049" md5="b707670ae423d0fda9fdb6f33e8f186c"
                     chunkPositionInArray="[0,0]" offset="130440821"
                     uuid="b0abe13e-4aab-47b3-b256-89d43380600e"/>
    </h4:chunks>
  </Float32>
</Dataset>
```

Use Cases

- Trade Study Use Case 1 DAP2 with CF
- Trade Study Use Case 2 DAP4 with CF
- Trade Study Use Case 3 DAP2 with the default option
- Trade Study Use Case 4 DAP4 with the default option
- Trade Study Use Case 5 FilenetCDF output
- Trade Study Use Case 6 Simulate NcML JoinExisting aggregation
- Trady Study Use Case 7 Simulate NcML JoinNew aggregation

Use case descriptions are available at:
https://github.com/OPENDAP/cloudydap/wiki

Usage and Cost Information

- Hourly AWS Cost and Usage Reports
- S3 access logs

Cost Modelling

- Fixed costs
  - EC2 instances (Hyrax servers)
  - Data in S3
- Dynamic costs
  - Number and type of user DAP requests
  - Number of Hyrax requests to S3
  - Outbound data
  - Cache type and size

Architecture #1 Cost Model

- One Hyrax server (EC2 instance)
- Original (unchanged) data files in S3
- Cache (EBS or EFS) and its size
- One S3 request for one data file

Architecture #2 Cost Model

- One Hyrax server (EC2 instance)
- Original (unchanged) data files in S3
- Cache is not necessary
- One or more S3 requests for one data file

Architecture #3 Cost Model

- One Hyrax server (EC2 instance)
- Only HDF5 datasets from original data files stored in S3
- Cache is not necessary
- One or more S3 requests per one HDF5 dataset

Acknowledgements

This work was supported by NASA contract number NNG15HZ39C.

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NASA, OPeNDAP Inc., Raytheon or The HDF Group.

OPeNDAP

The HDF Group

Raytheon