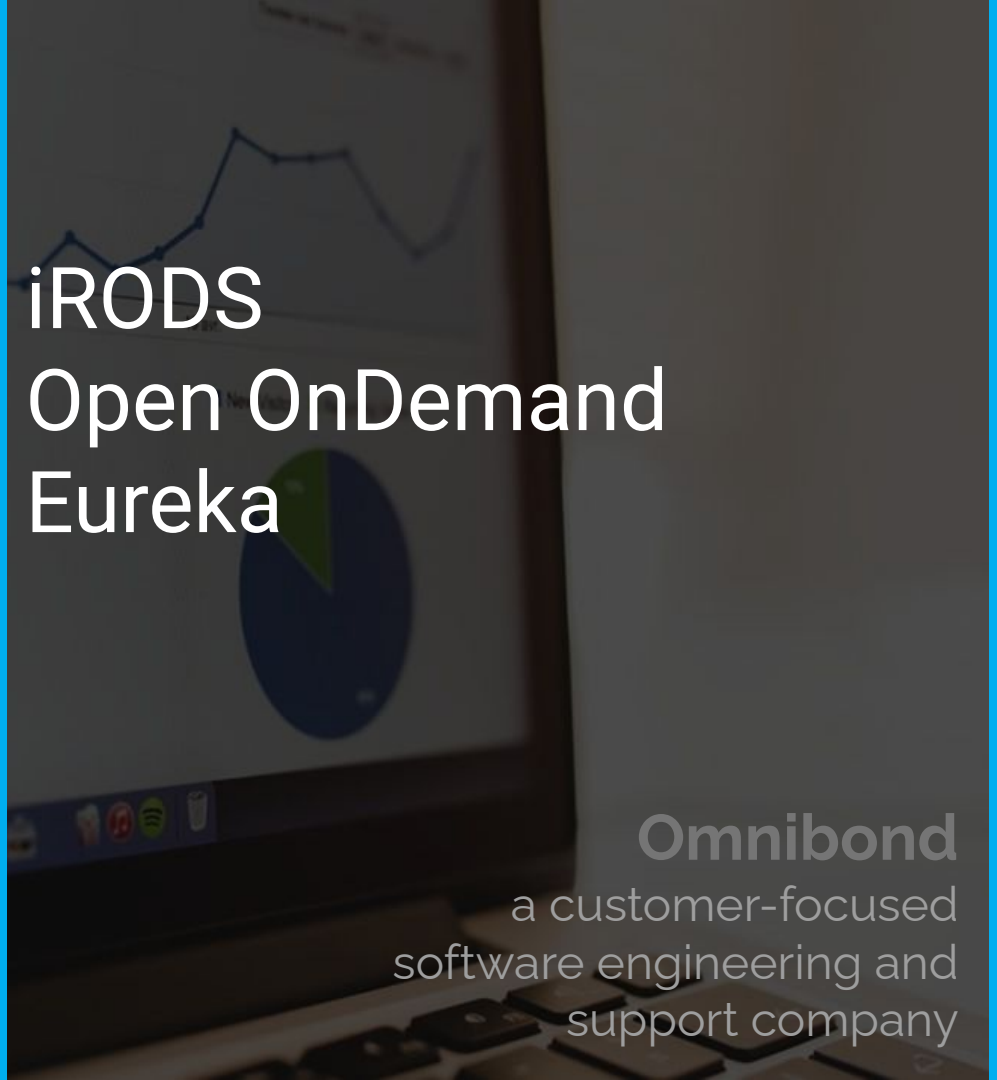

iRODS
Open OnDemand
Eureka

Omnibond

a customer-focused
software engineering and
support company

Leadership Team

Omnibond

a customer-focused software engineering and support company

- Over 40 years combined experience in facilitating and supporting academic and corporate research in using the tools and technologies of advanced computing
- Experience at the working and technical, project and executive management levels at Clemson, Purdue and Miami Universities and the National Center for Supercomputing Applications located at UIUC
- Expertise in software development, systems integration, operations, applications support, data transmission, identity and access management, customer relations and research facilitation and engagement
- Founded ACI-REF http://www.aciref.org and CaRCC http://carcc.org
- Over three decades of funded projects from NSF, DoD, DoE, NSA, NIST and DARPA
- INTERNET2 Presidential Fellow & CSTAAC Committee Members

- Identity & Security Management
  - NetIQ Identity Manager Connectors
  - Thousands of customers, sold through Novell/Micro Focus/OpenText, since early 2000’s
  - OmniPasskey Passwordless MFA Shibboleth Plug-in
- Computer Vision & AI
  - TrafficVision - AI based Automated Incident Detection (AID) & Data from existing cameras on roadways
  - BayTracker - Retail Vehicle Tracking and Timing
  - Port Observer - Drayage Queuing, AIS, Dashboard for Ports
- Cloud HPC and Storage Orchestration
  - CloudyCluster
  - OrangeFS
  - Eureka Project
  - Custom Cloud <-> On-Prem Integration

Software Products

Omnibond

a customer-focused software engineering and support company

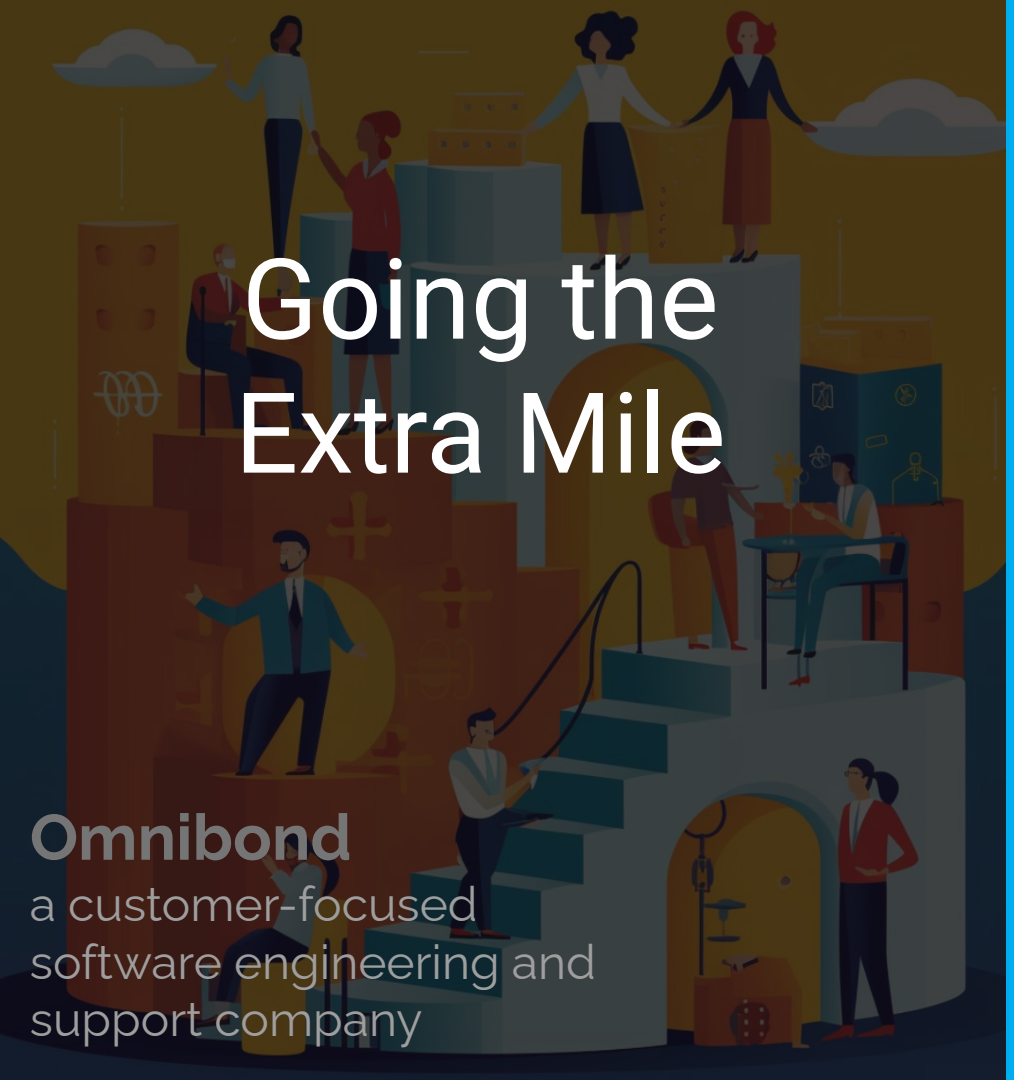

I have worked with Omnibond development and support for almost 15 years. They always deliver the highest level of support possible. I have been in Technical Support for 23 years and have worked with different companies across the world. Omnibond is at the top in response time and for going the extra mile to help our customers. I have no reservation in recommending them as a great company.

Reed Harrison

Provo, UT

Support Engineer

Novell/NetIQ/MicroFocus/Opentext

Real-time HPC & AI TrafficVision

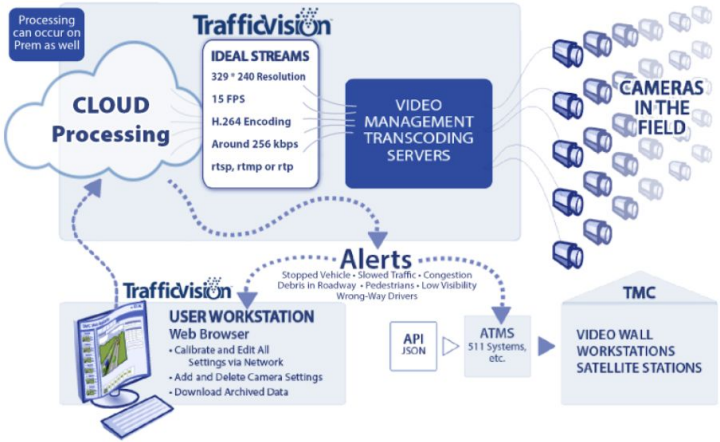

Features

Incident Detection

- Stopped Vehicle
- Slowed Traffic
- Debris in Roadway
- Low Visibility
- Pedestrians
- Wrong-Way Drivers*

Alerts, Incident Images & Clips

Real-Time Data Collection

- Speeds Per Lane / Per Direction
- Vehicle Counts Per Lane / Direction
- Lane Occupancy and Density
- Congestion Index
- Classification:
  **Motorcycle / Car / Truck / Large Truck

Data CSV export & JSON API

Customers

Caltrans · CDOT · Calgary · FDOT · 360ns · GDOT · SCOUT · Skyline Technology Solutions · kapsch · WSP · Ontario · Michigan DOT · Minnesota Department of Transportation · MoDOT · NMDOT · NTTA · Pennsylvania Turnpike · Peraton · Port of Long Beach · RCTC · SCDOT · TRIMARC

Scaling on AWS

AWS News Blog

Natural Language Processing at Clemson University – 1.1 Million vCPUs & EC2 Spot Instances

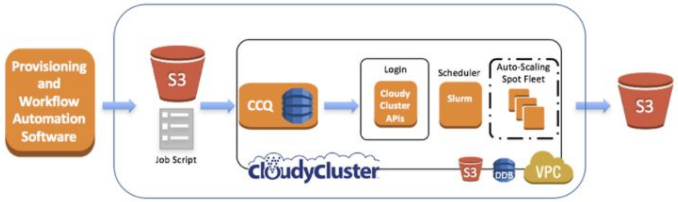

Provisioning and Workflow Automation Framework

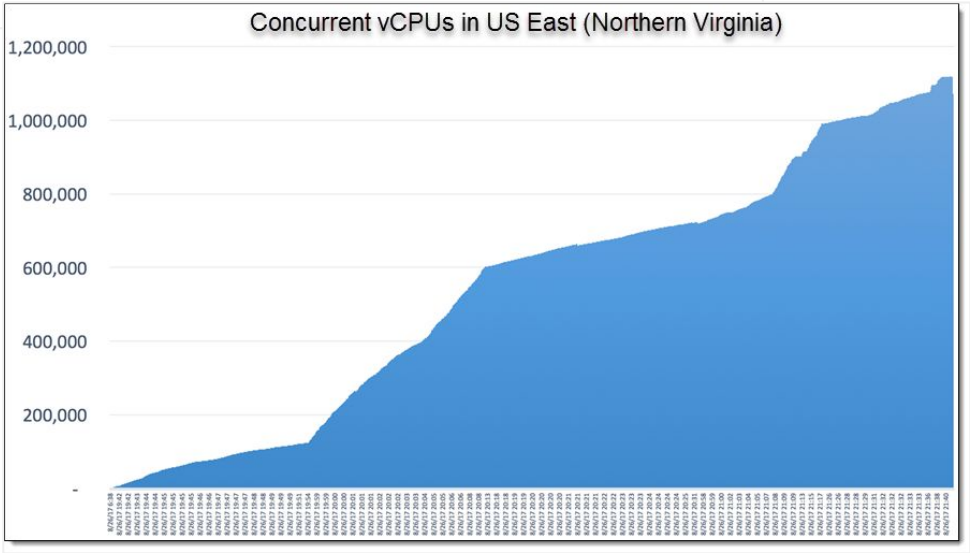

Data Intensive Computing Ecosystems (DICE)
School of Computing, Clemson University

https://aws.amazon.com/blogs/aws/natural-language-processing-at-clemson-university-1-1-million-vcpus-ec2-spot-instances/

Scaling

Google HPC Blog Post

Cloud against the storm: Clemson's 2.1 million VCPU experiment

https://cloud.google.com/blog/topics/hpc/clemson-experiment-uses-2-1-million-vcpus-on-google-cloud

THENEXTPLATFORM

Urgent HPC can Burst Affordably to the Cloud

https://www.nextplatform.com/2020/01/08/urgent-hpc-can-burst-affordably-to-the-cloud/

- 133,573 GCP Instances at peak
- 2,138,000 vCPUs at peak
- 6,022,964 vCPU hours

Processed 2,479,396 hours (~256TB) of video data

- ~4 hours of runtime
- ~1M vCPU within an hour
- ~1.5M vCPU within 1.5 hours
- 2.13M vCPU within 3 hours

Total Cost: $52,598.64 USD
Average cost of $0.008 USD per vCPU hour

CLEMSON UNIVERSITY

Google Cloud

intel

Cloudy Cluster by Omnibond

TrafficVision

Turn-Key Elastic HPC, HTC, & Storage in AWS

Create a familiar, secure & fully operational computational cluster in minutes, complete with:

Encrypted Storage: EBS, OrangeFS on PD

Compute: Job Driven Elastic Compute through CCQ (Spot, On-Demand, GPU)

Schedulers: Torque & SLURM with the CCQ Meta-Scheduler
Supports Billing Tags per job/user
Placement Policies, Advanced Networking
Supports Multiple Custom Images

End User Friendly UI: Open OnDemand

Includes Familiar Software

HPC Libraries: Boost, Cuda Toolkit, Docker, FFTW, FLTK, GCC, Gengetopt, GRIB2, GSL, HDF5, Intel MPI, Intel Runtimes, ImageMagick, JasPer, mpich, NetCDF, NumPy, Octave, OpenCV, OpenMPI, PROJ, R, Rmpi, SciPy, SWIG, WGRIB, UDUNITS, .NET Core, Singularity, Queue, Picard, xrootd, etc…

HPC Software: Ambertools, ANN, ATLAS, BLAS, Blast, Blender, Burrows-Wheeler Aligner, CESM, GROMACS, JupyterLab, LAMMPS, NCAR, NCL, NCO, nwchem, OpenFoam, papi, paraview, Quantum Espresso, SAMtools, WRF, Galaxy, Vtk, Su2, Dakota, Gatk, etc…

ML Software: Mlpack, NuPIC, Octave, OpenCV, PICARD, Queue, Scikit-learn, Tensorflow

Sample Jobs: Intel Cluster Checker, Mpi_prime, wrf, etc…

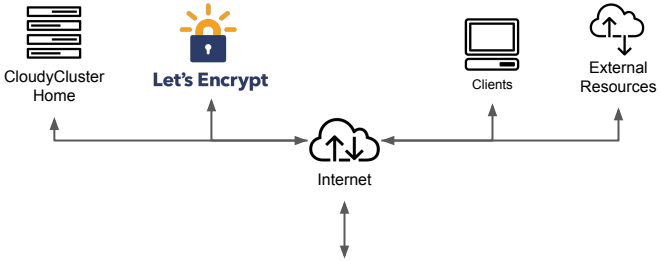

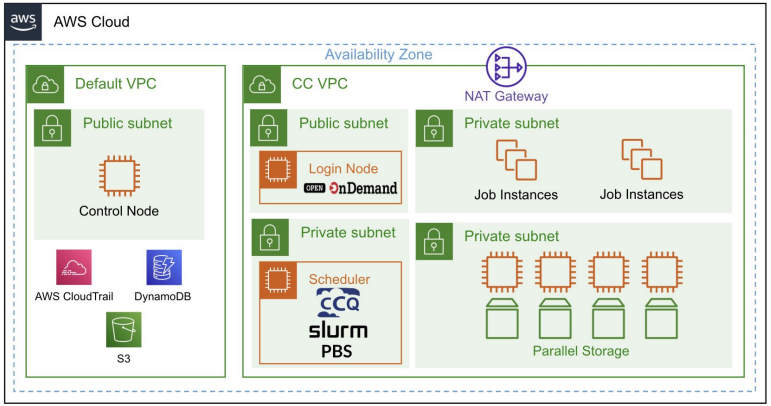

cloudycluster

Turn-Key Elastic HPC, HTC, & Storage in GCP

Create a familiar, secure & fully operational computational cluster in minutes, complete with:

Encrypted Storage: GCS, OrangeFS on PD

Compute: Job Driven Elastic Compute through CCQ (Preemptable, On-Demand, GPU)

Schedulers: Torque & SLURM with the CCQ Meta-Scheduler
Supports Billing Labels per job/user
Placement Groups, BulkAPI, Tier1 Networking
Supports Multiple Custom Images

End User Friendly UI: Open OnDemand

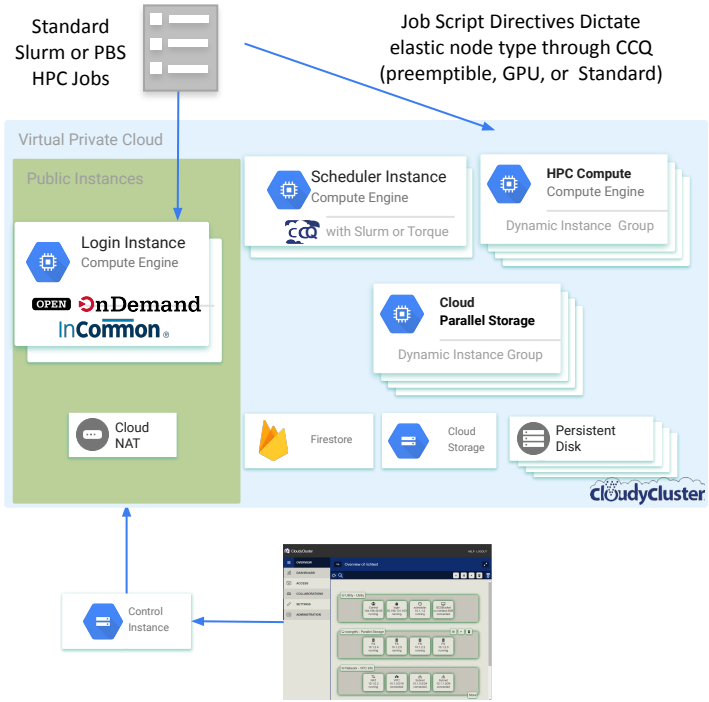

Includes Familiar Software

HPC Libraries: Boost, Cuda Toolkit, Docker, FFTW, FLTK, GCC, Gengetopt, GRIB2, GSL, HDF5, Intel MPI, Intel Runtimes, ImageMagick, JasPer, mpich, NetCDF, NumPy, Octave, OpenCV, OpenMPI, PROJ, R, Rmpi, SciPy, SWIG, WGRIB, UDUNITS, .NET Core, Singularity, Queue, Picard, xrootd, etc…

HPC Software: Ambertools, ANN, ATLAS, BLAS, Blast, Blender, Burrows-Wheeler Aligner, CESM, GROMACS, JupyterLab, LAMMPS, NCAR, NCL, NCO, nwchem, OpenFoam, papi, paraview, Quantum Espresso, SAMtools, WRF, Galaxy, Vtk, Su2, Dakota, Gatk, etc…

ML Software: MIpack, NuPIC, Octave, OpenCV, PICARD, Queue, Scikit-learn, Tensorflow

Sample Jobs: Intel Cluster Checker, Mpi_prime, wrf, etc…

cloudycluster

Google Cloud Marketplace

HPC User Experience

OPEN OnDemand

Integration of OOD into CloudyCluster

- As part of the Cloudify Gateways program (Omnibond, OSC, U Buffalo, Va Tech)
- Full deployment automated as part of the CloudyCluster launch
- The Collaboration Continues

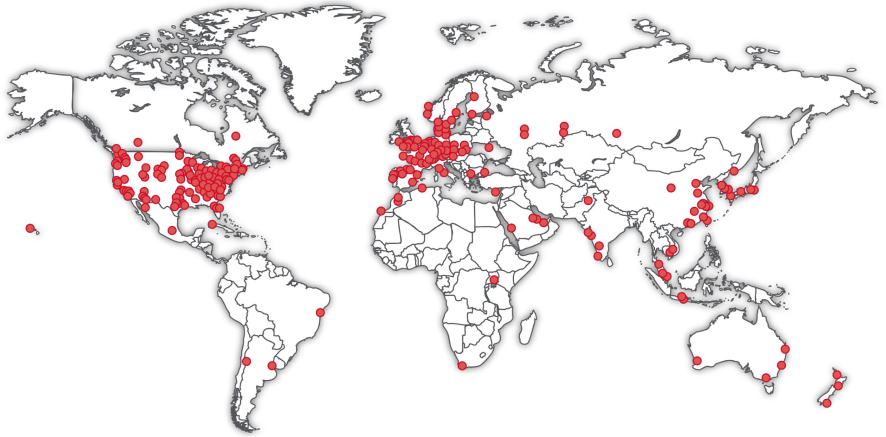

Open OnDemand Deployments

"Your browser is the supercomputer: On Demand is a no-tears shortcut to research-computing" from Matt Windsor of University of Alabama at Birmingham. A key phrase: "No experience necessary"

This material is based upon work supported by the National Science Foundation under grant numbers 1534949 and 1835725. and under active development by a team from the Ohio Supercomputer Center, U. of Buffalo CCR, and Virginia Tech

Relion & CRYO-EM

Working with the Cianfrocco Lab at University of Michigan to test data transfer and scalability to the Cloud for CRYO-EM projects using RELION

Also working with related Startups in the space that don't have access to large on prem clusters.

Initially Introduced as part of the Cloudify Program through the SGCI.

Cianfrocco is willing to consult with customers in conjunction with Omnibond and CloudyCluster

Another CRYO-EM project is in progress with the ERN ECOSYSTEM for RESEARCH NETWORKING

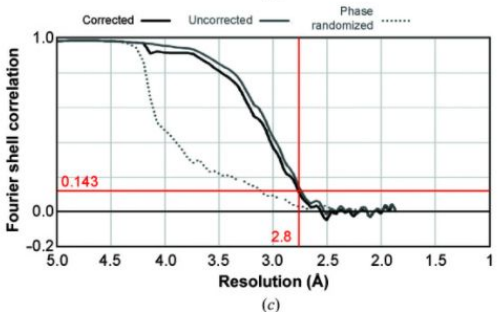

High-resolution cryo-EM using beam-image shift at 200 keV

Jennifer N Cash [1], Sarah Kearns [1], Yilai Li [1], Michael A Cianfrocco [1]

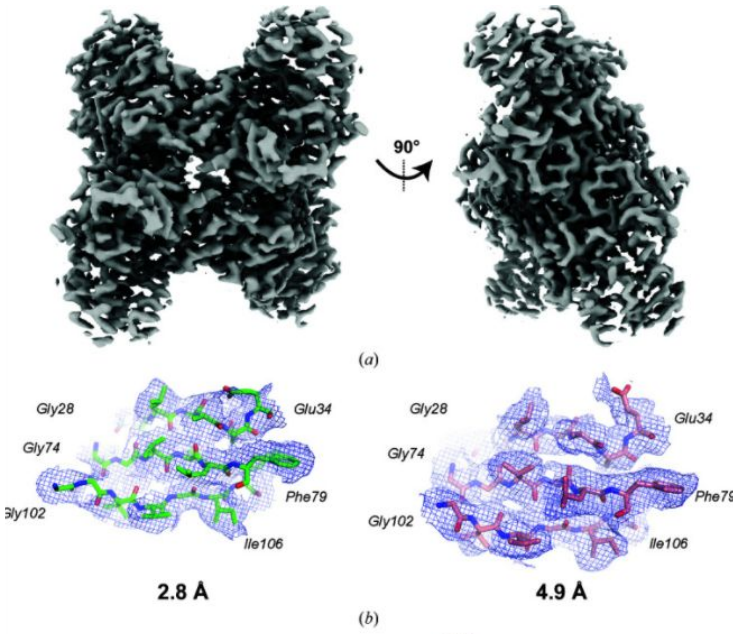

Figure 5 Final aldolase reconstruction at 2.8 Å resolution. (*a*) Sharpened aldolase reconstruction at 2.8 Å resolution. (*b*) Example densities and models for aldolase at 2.8 and 4.9 Å resolution. (*c*) FSC curve for the final reconstruction.

SDPB Solver MPI Tuning

The Problem of Strong Coupling

Quantum field theory (QFT) is a universal language for theoretical physics, describing the Standard Model of particle physics, early universe inflation, and condensed matter phenomena such as phase transitions, superconductors, and quantum Hall fluids. A triumph of 20th century physics was to understand weakly coupled QFTs. However, weakly interacting systems represent a tiny island in theory space and cannot capture many of the most interesting physical phenomena.

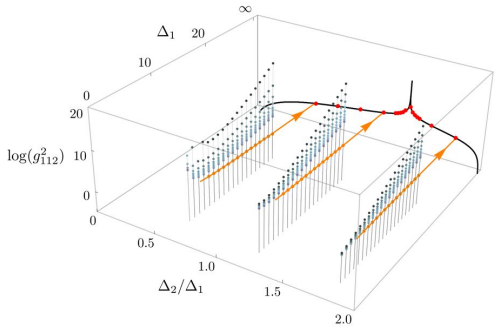

https://bootstrapcollaboration.com/

Our Collaboration

Recently, members of our collaboration discovered new bootstrap techniques that apply in general dimensions. In the past few years we have applied these techniques to a wide variety of seemingly unrelated problems: to perform the world's most precise analysis of the 3d Ising model, to constrain strongly coupled theories of physics beyond the Standard Model, to aid in classifying superconformal field theories, to derive locality and black hole thermality in models of quantum gravity, and to prove irreversibility of renormalization group flows. We believe this is the beginning of a much larger enterprise, crossing traditional boundaries between string theory, condensed matter physics, and phenomenology, and making strong connections to modern mathematics and computer science.

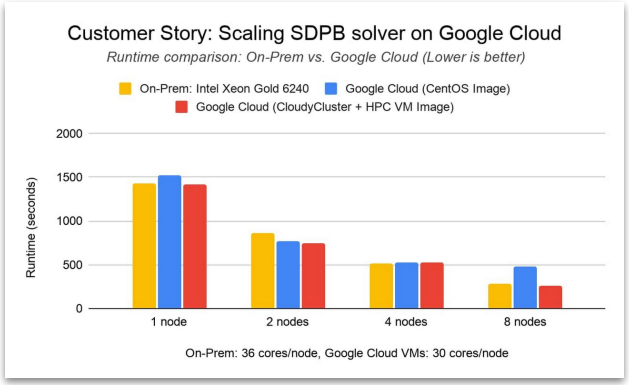

https://cloud.google.com/blog/topics/hpc/introducing-hpc-vm-images

To expand the collaboration's computation capabilities, Walter Landry wanted to see how SDPB would scale on Google Cloud. Working with Omnibond's **CloudyCluster** and leveraging the HPC VM image, Landry achieved comparable performance and scaling to an on-premises cluster at Yale, based on Intel Xeon Gold 6240 processors and Infiniband FDR.

TAPIS - PEARC ‘22 Poster, Gateways ‘22 Paper

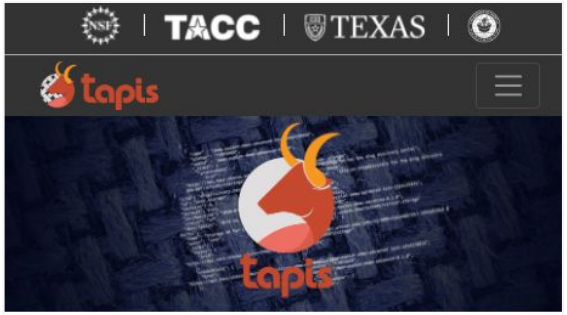

Tapis

Unify, simplify, automate and secure your interactions with advanced computing infrastructure.

The Tapis Framework provides a hosted, unified web-based API for securely managing computational workloads across institutions so that experts can focus on their research instead of the technology needed to accomplish it.

Projects Using Tapis

- A2CPS
- Bridging Barriers
- CyVerse
- DesignSafe
- Drug Discovery Portal
- ECCO
- Ike ‘Wai (at the University of Hawaii Manoa)
- iReceptor
- iR+
- NeuroNex 3DEM
- Planet Texas 2050
- Science Gateways Community Institute
- Synergistic Discovery and Design Environment (SD2E)
- VDJServer
- UTRC

https://github.com/omnibond/automaton

Extending Tapis Workflow Management Framework with Elastic Google Cloud Distributed System using CloudyCluster by Omnibond

ERIC LAM, University of Hawaii, Manoa, USA
SEAN CLEVELAND, University of Hawaii - Information Technology Services - Cyberinfrastructure, USA
COLE MCKNIGHT, Omnibond Systems, USA
BOYD WILSON, Omnibond Systems, USA
RICHARD CARDONE, Texas Advanced Computing Center, USA
JOE STUBBS, Texas Advanced Computing Center, USA

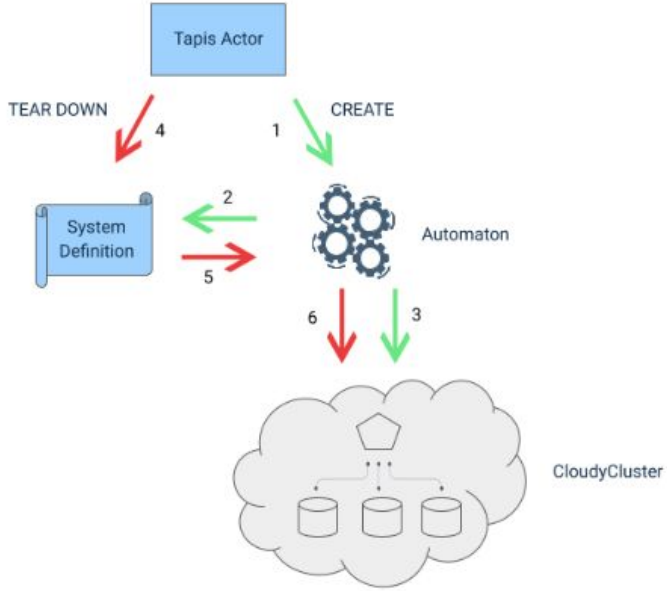

The Meta-Scheduler Approach

Scheduler Independent

- Torque
- Slurm

Allow for Meta-Scheduler Directives

- Instance Type, each job can have a different instance type and billing model.
- GPU
- Spot, Preemptible
- Billing Tags/Labels
- Volume Type
- Custom Images for Different Jobs

Turns “Scheduler” into Dispatcher Handling:

- Instance and appropriate subnet provisioning
- Instance deletion

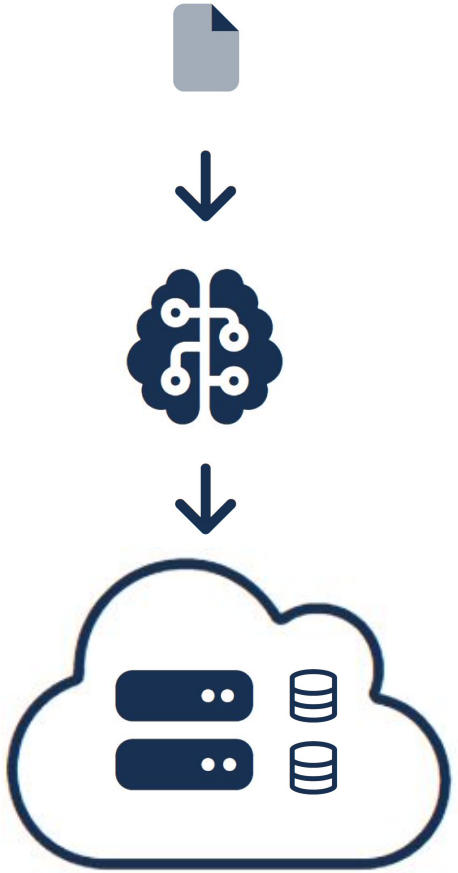

Project

EUREKA

see the show

OOD UI + Eureka Vision

- Interactive
  - Applications & Launchers
  - API Applets & Saas Apps
  - Project Focused
- Compute
  - Compute Anywhere (HPC, AI, & Beyond)
  - Enable Cloud Specialties
  - Simplify Compute and Storage Interactions
- Storage
  - Integrate Diverse Storage Resources
  - Collaborate First
  - Project Level Data Lifecycle

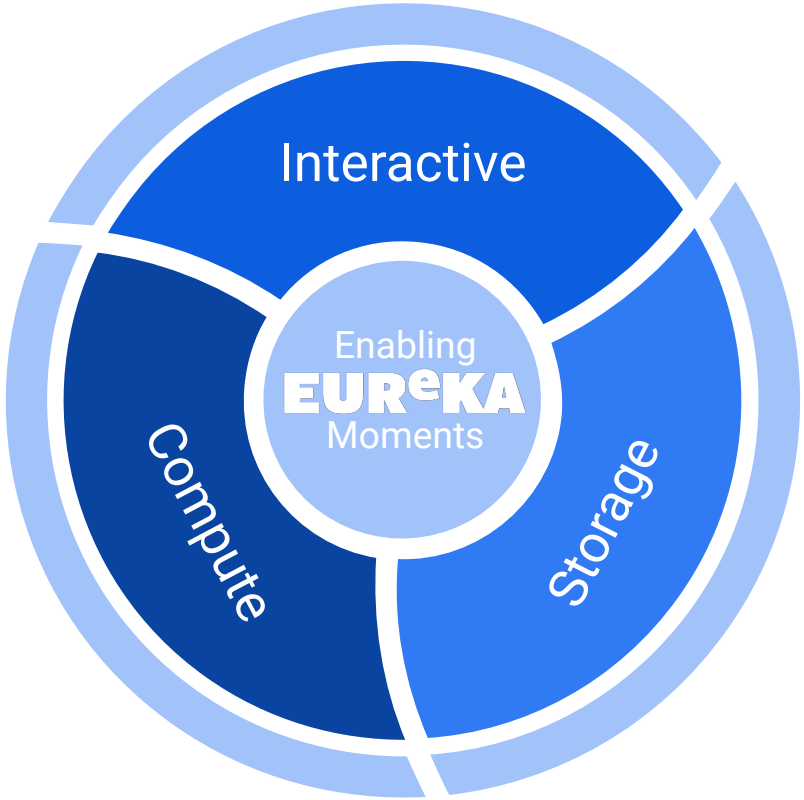

Technology

- **Cloud Jump UI & Open OnDemand**
  - HPC Industry Standard
  - OODUI - Simplify & Empower
- **Omni-Scheduler**
  - Coordinates UI, Storage, & Compute
- **Constellation Driven Terraform**
  - Multi-Cloud Enabler
- **iRODS**
  - Metadata Driven Data Management & Movement
  - Designed for Diverse Storage

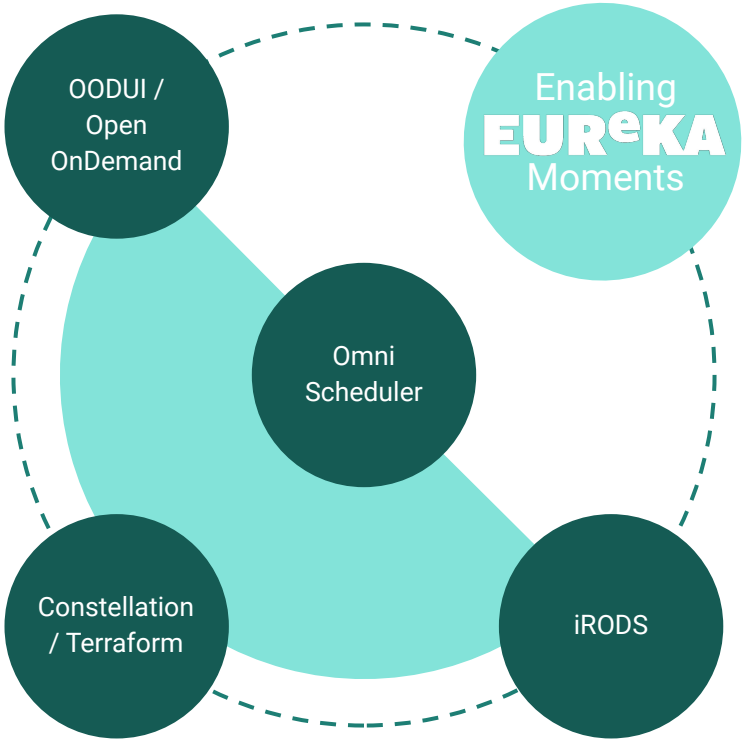

Inside-Out

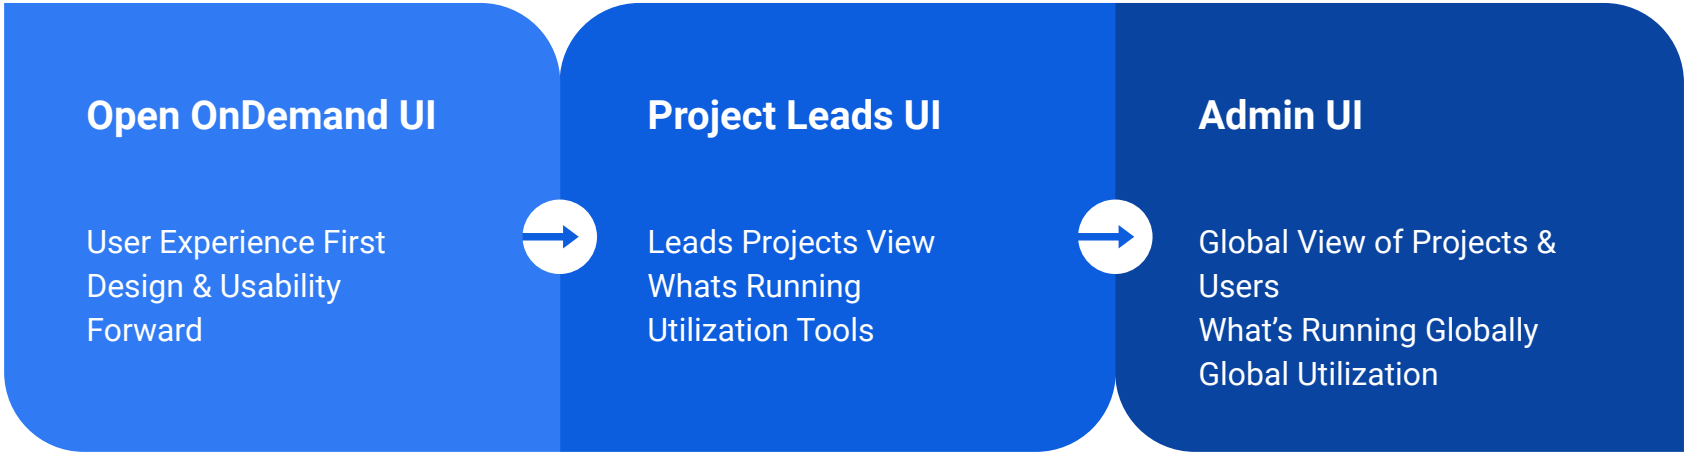
Open OnDemand UI
User Experience First
Design & Usability
Forward
Project Leads UI
Leads Projects View
Whats Running
Utilization Tools
Admin UI
Global View of Projects &
Users
What's Running Globally
Global Utilization

Open OnDemand Example Deployments

Don't see an organization?

The more the merrier! Let us know any that belong on the list

openondemand.org/orgs

Eureka - Architecture (Work in Progress) 5/2023

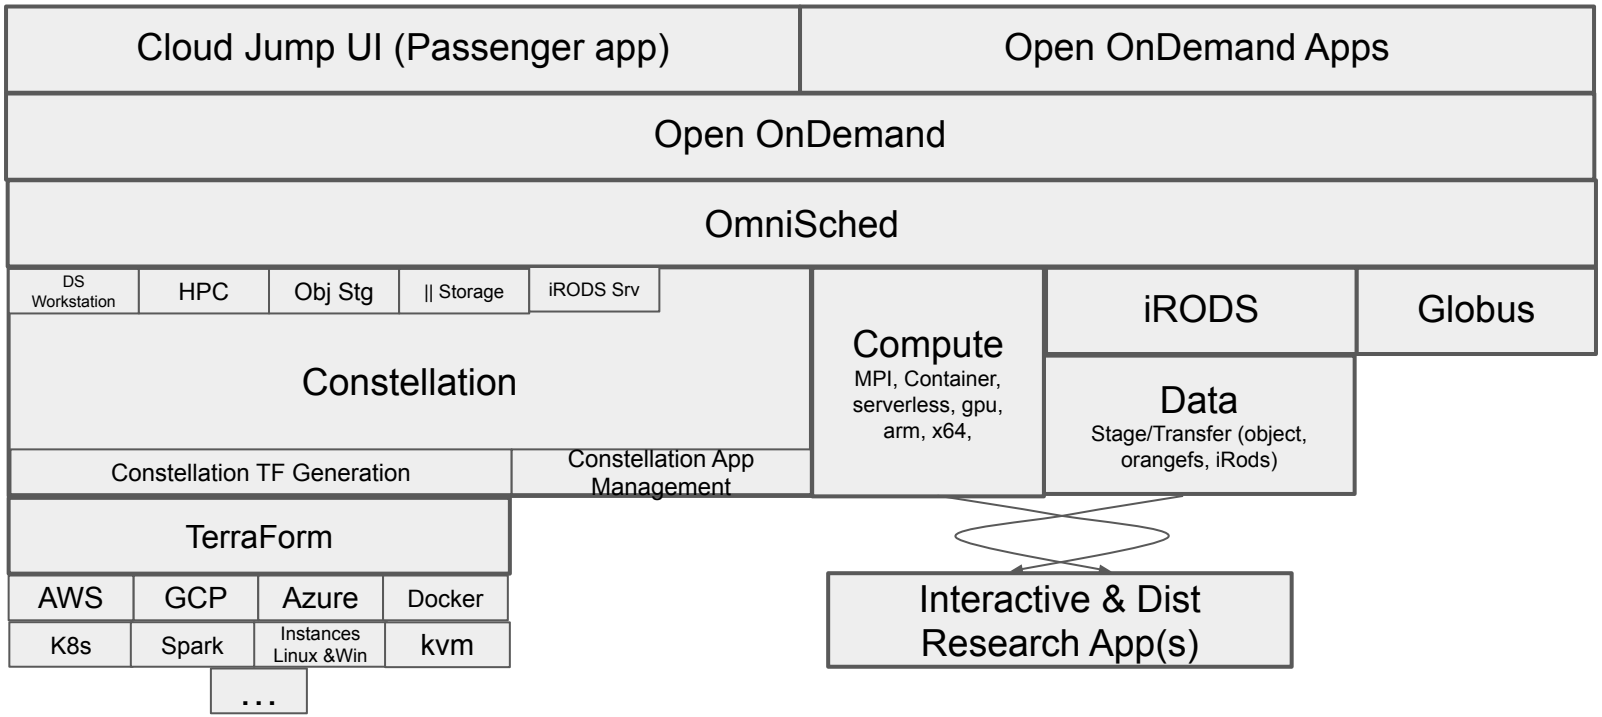

The Omni-Scheduler Approach

Expand meta-scheduler concept (OmniSched) to handle:

- Cross Cloud - leveraging Constellation / TerraForm
  - Enable On-Prem (TF -> kvm)
- Job routing between locations
- Dynamic Scratch
  - Driven by job Directives
- Data Staging
  - Driven by Job Directives (Input Data & Results)
- Beyond HPC
  - Data Science Apps
  - HPC adjacent (publishing, websites, etc..)
  - Enable Apache Spark
  - K8s
  - Etc…
- Scheduling Data Movements
  - iRODS integration

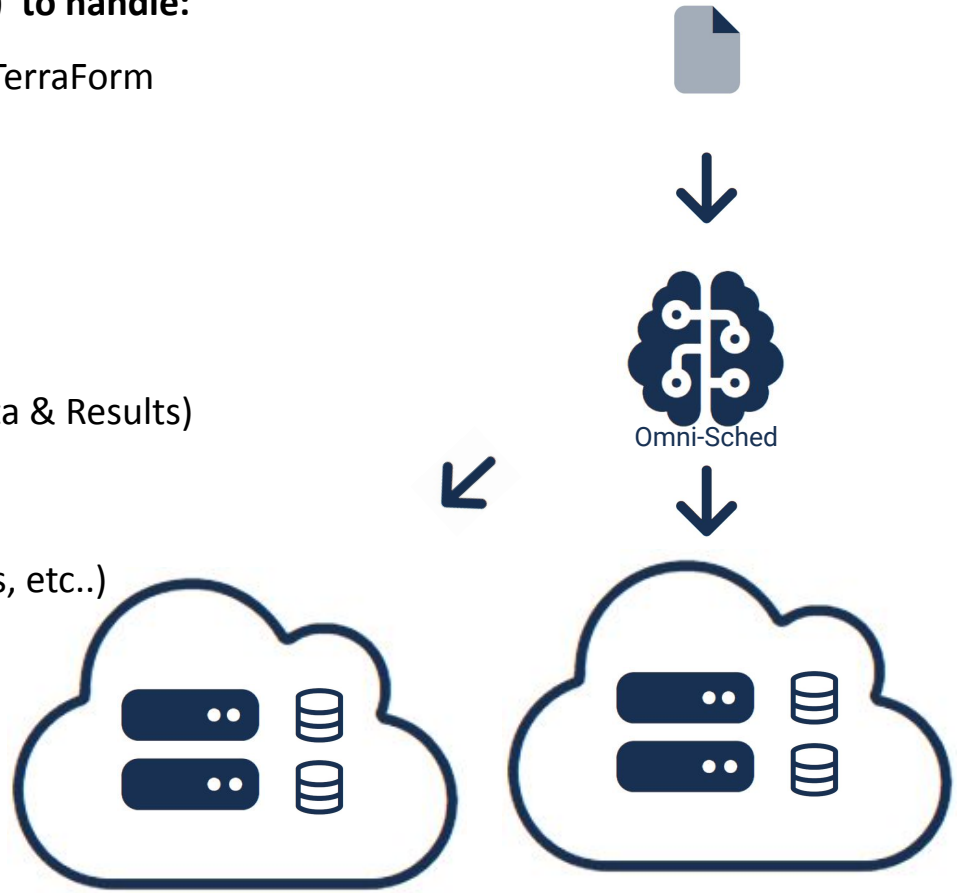

The Integrated Data Management

In addition to Job based directives

Data Management Capabilities

- Replicate
- Archive
- Ingest
- Tier
- Publish

All configured in OOD with iRODS

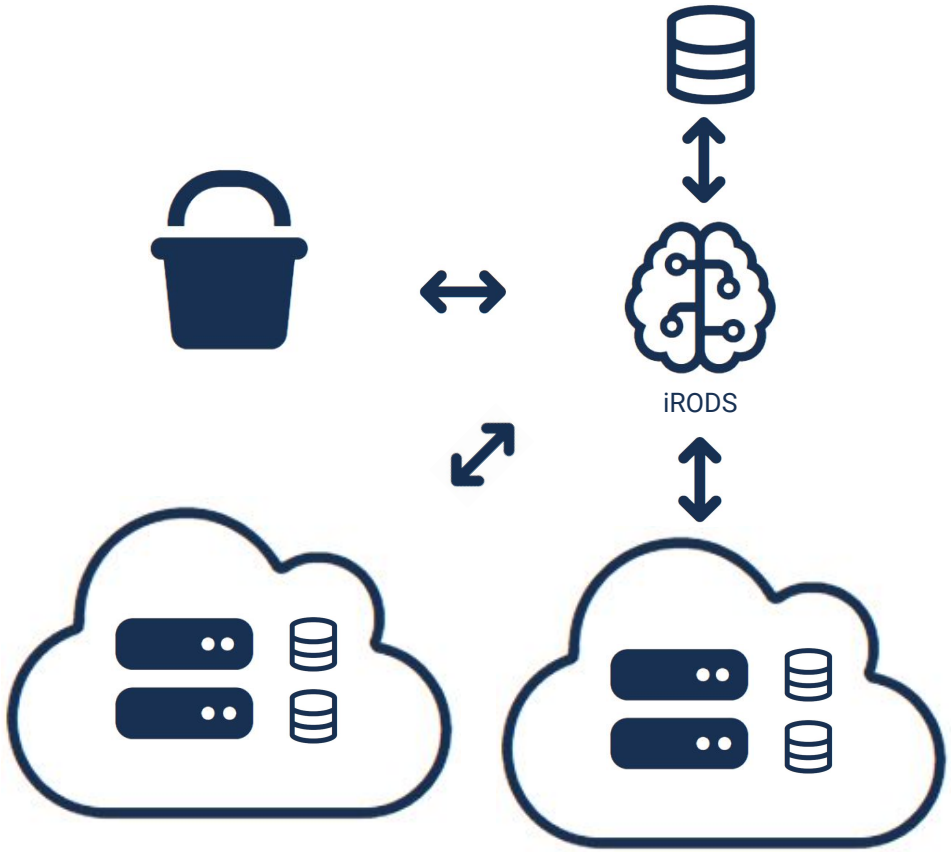

Current Deployment Architecture 5/2023

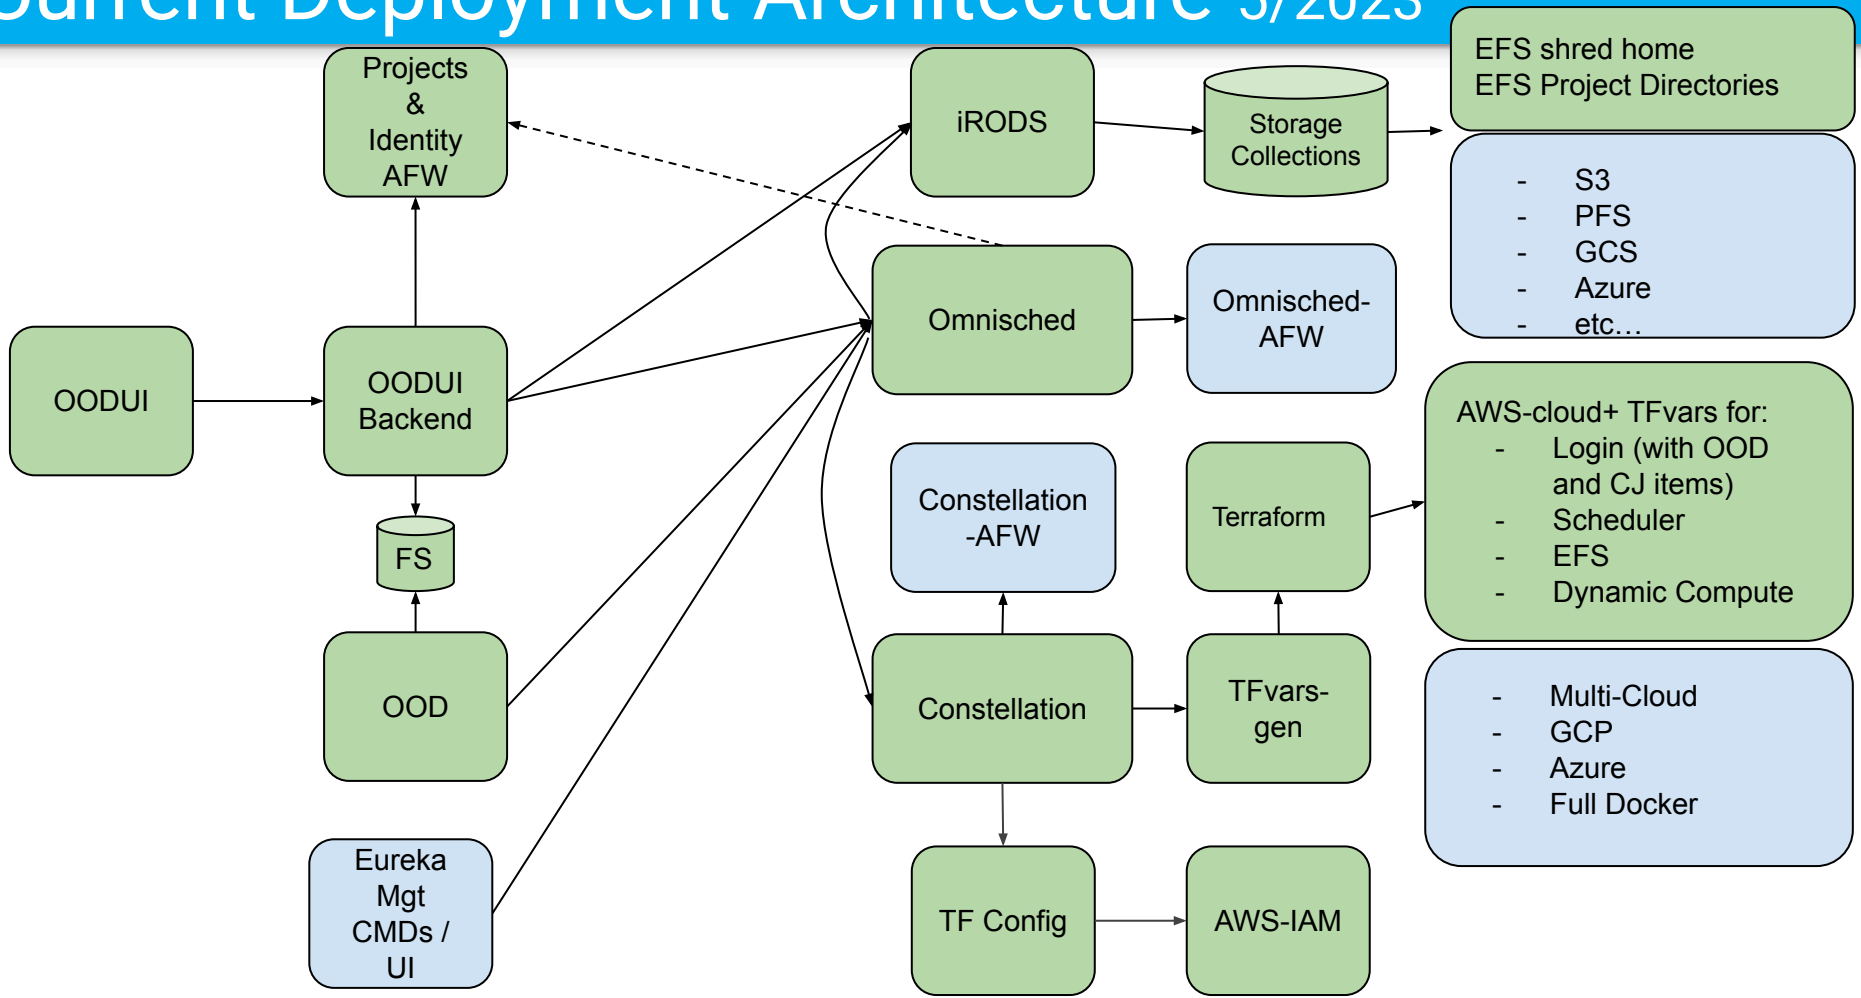

Deployment Architecture

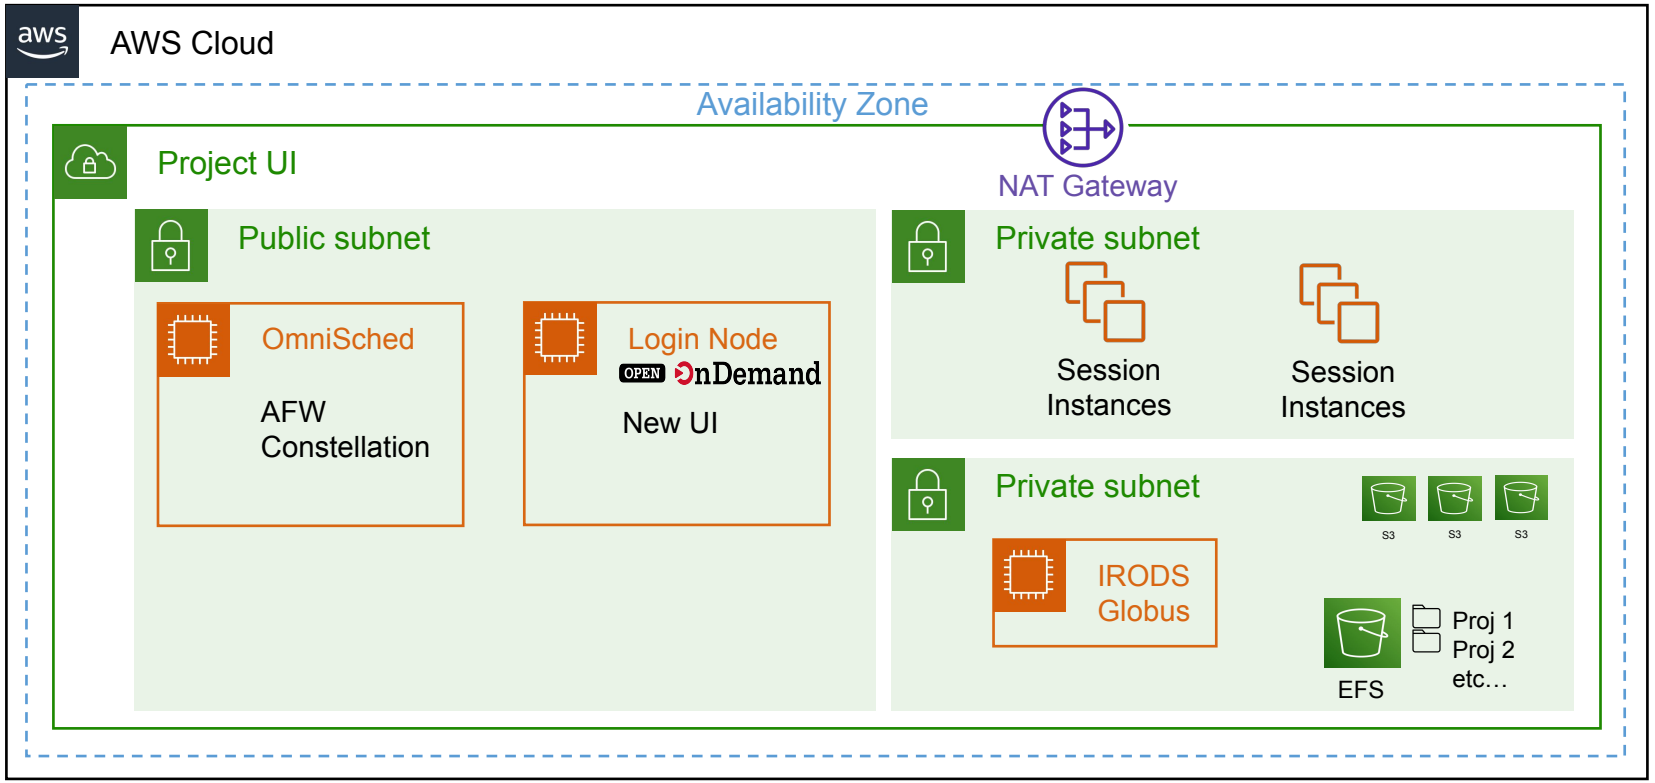

HBS CloudJump

Destination

EUREKA

Data Science, AI, & HPC

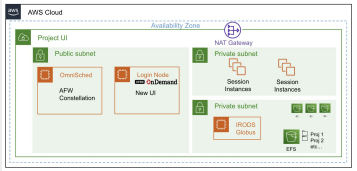

GCP

On-Prem

AWS

Docker, Azure, etc...

Demo

EUREKA

see the show

Cloud Jump

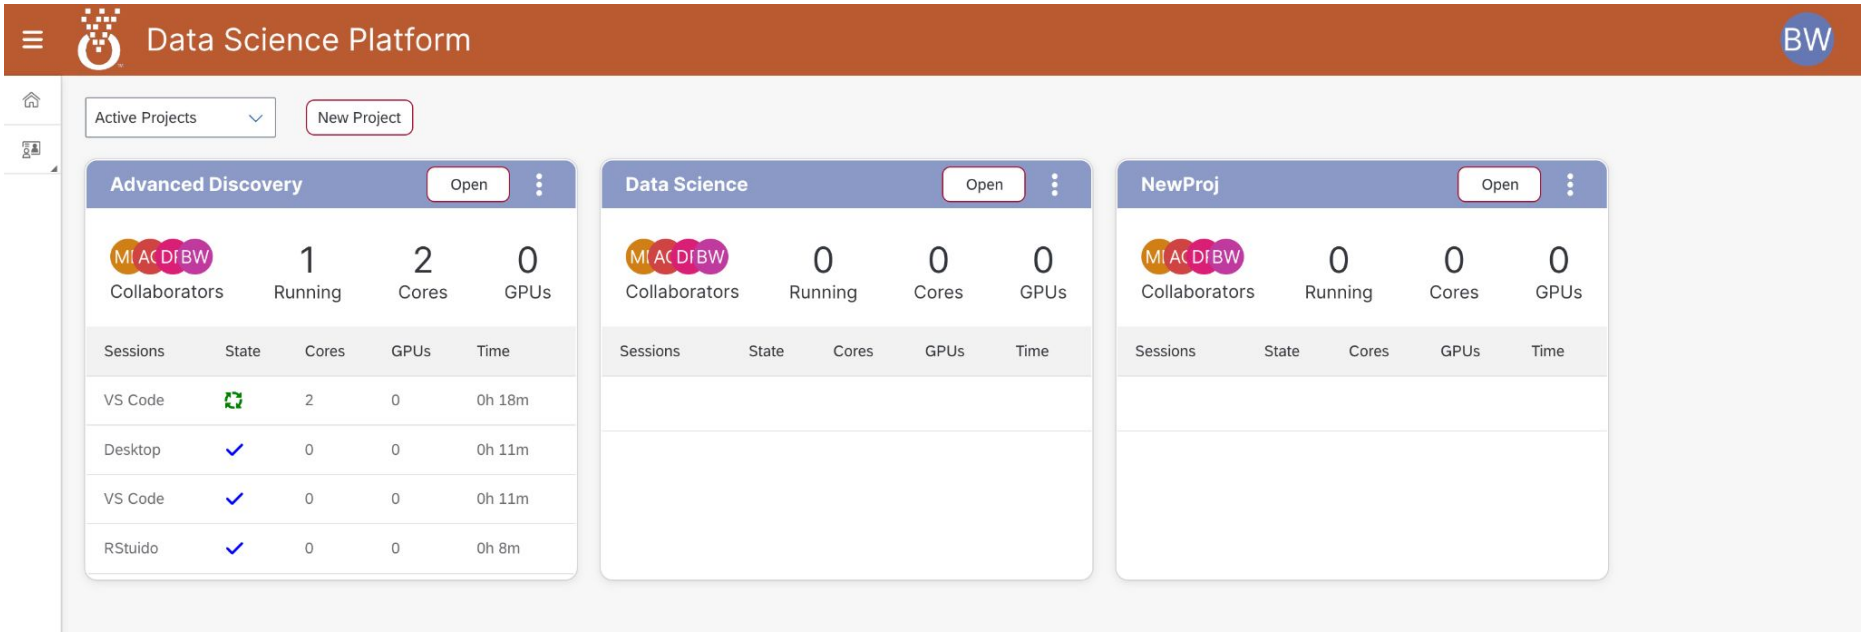
Data Science Platform
BW
Active Projects
New Project
Advanced Discovery
Open
Collaborators
1
Running
2
Cores
0
GPUs
Sessions
State
Cores
GPUs
Time
VS Code
2
0
0h 18m
Desktop
0
0
0h 11m
VS Code
0
0
0h 11m
RStuido
0
0
0h 8m
Data Science
Open
Collaborators
0
Running
0
Cores
0
GPUs
Sessions
State
Cores
GPUs
Time
NewProj
Open
Collaborators
0
Running
0
Cores
0
GPUs
Sessions
State
Cores
GPUs
Time

Project Eureka - Project Based UI

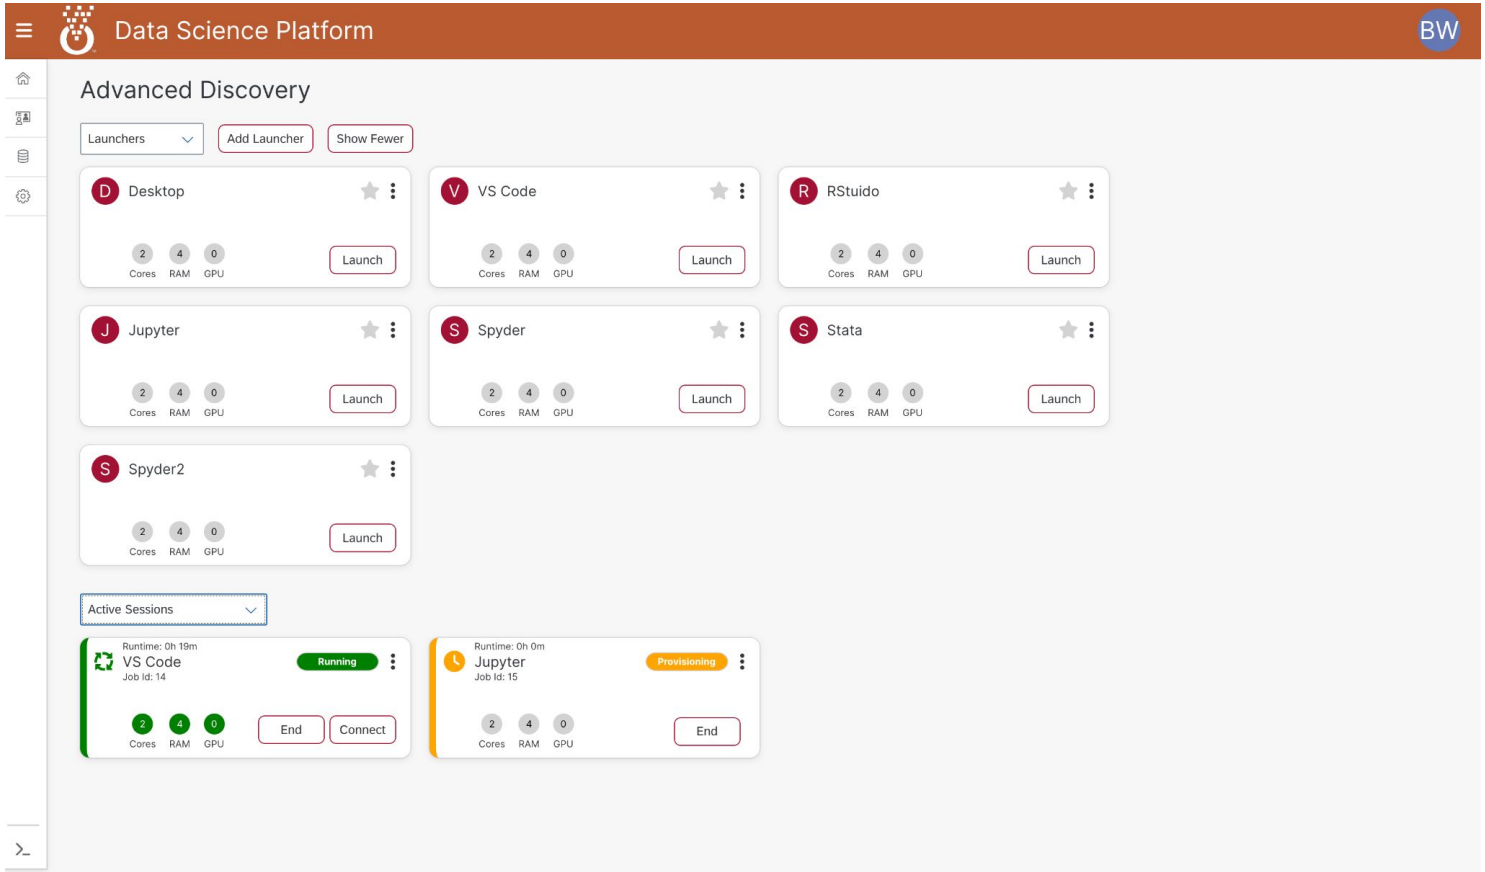

Future Storage Collections

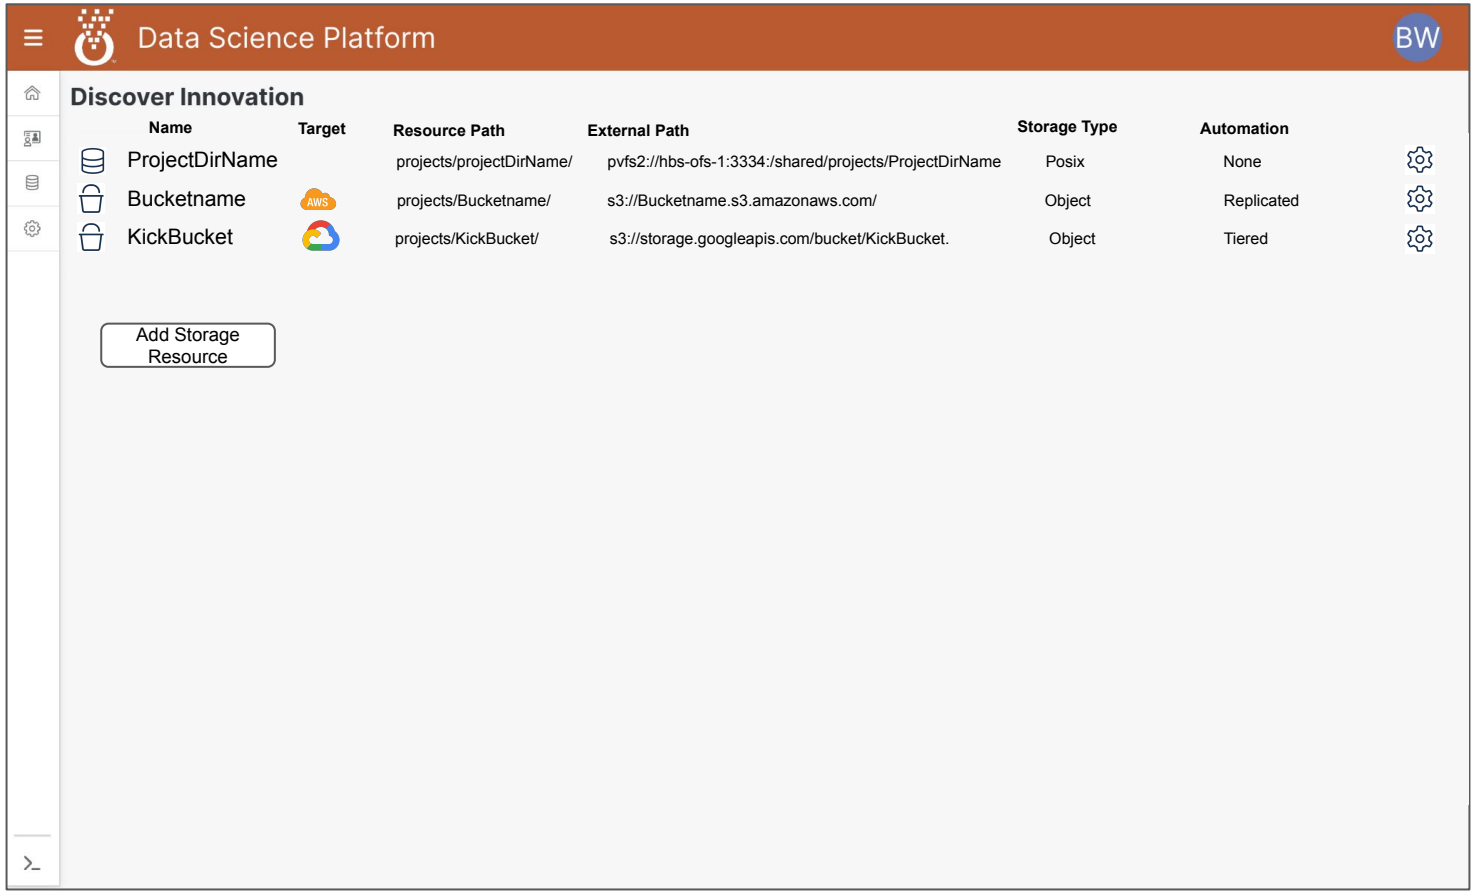

Future Collections Access in Storage Manager

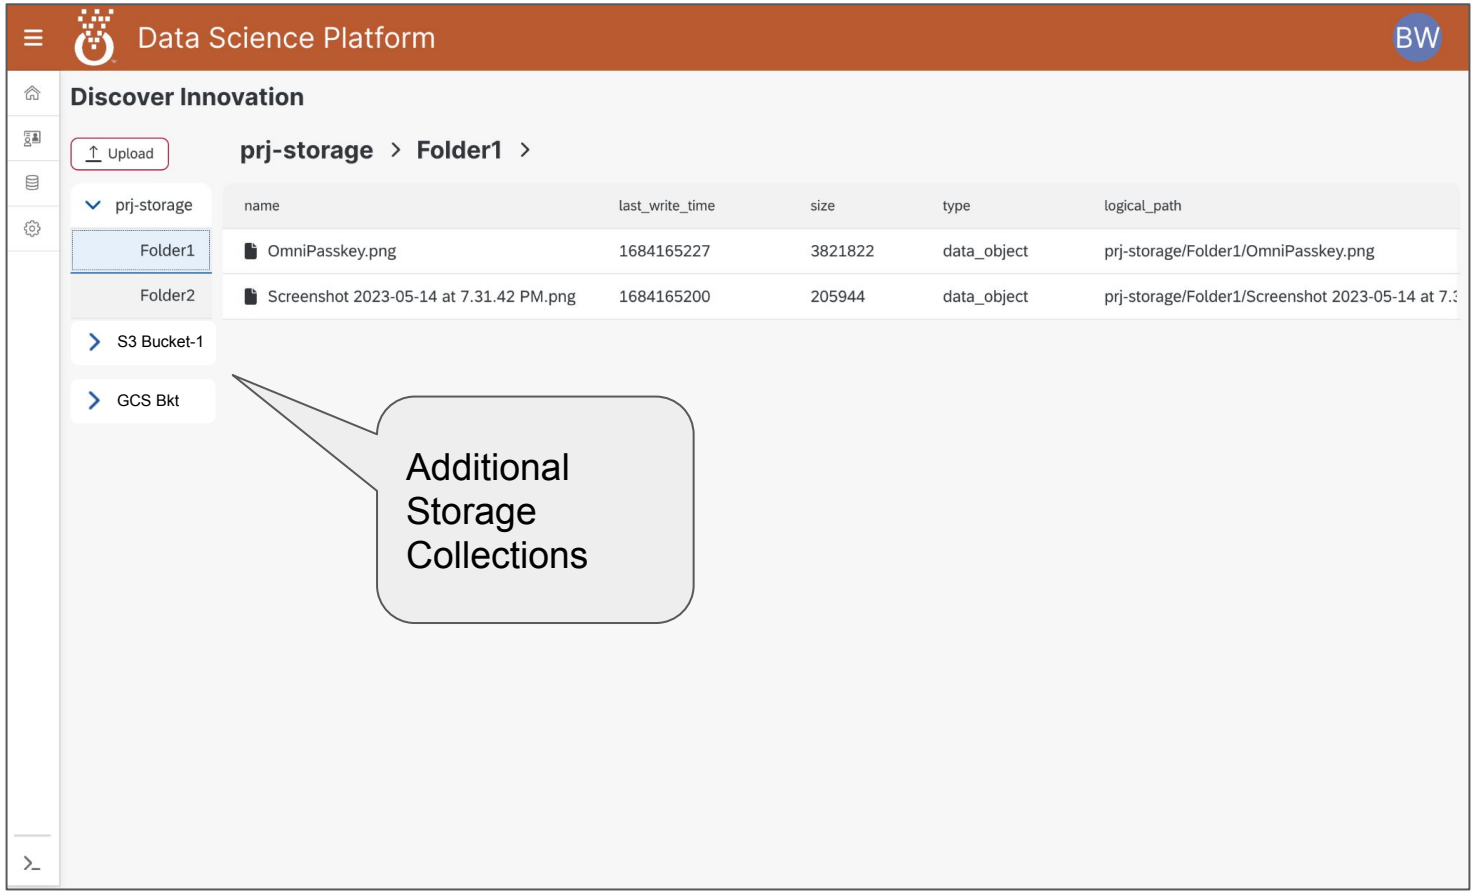

iRODS: What we're using (so far…)

- iCommands
  - Initialization (default collections, groups, permissions, and metadata)
- REST Client
  - Initialization:
    - Create corresponding iRODS user for each LDAP user
  - Frontend file manager:
    - List files/collections
    - Create collections/files
    - Rename files
    - Upload/download files (stream data)
- NFSRODS
  - Allow apps (e.g., VS Code Server) to directly access iRODS files
- iRODS Rule Language (NREP)
  - Automatically tag collections with custom metadata
- Next
  - OrangeFS MD Syncer (dynamic scratch and linux home sync to iRODS)
    - TBD: How to notify the catalog of file adds, deletes, updates

Tentative Timeline

Storage Manager - Download, Hidden Files, Multi-Select, create menu at top
Instances - Start Stop vcpu/ram display in list
Project - directory prefill,
Add Launcher - option in left menu
iRODS - s3 bucket create

Jobs - per job scratch, data staging, billing tags
Launcher - file search
AdminUI - Initial Admin UI
iRODS - OranegFS Syncer
Platform - local Docker Compute Nodes

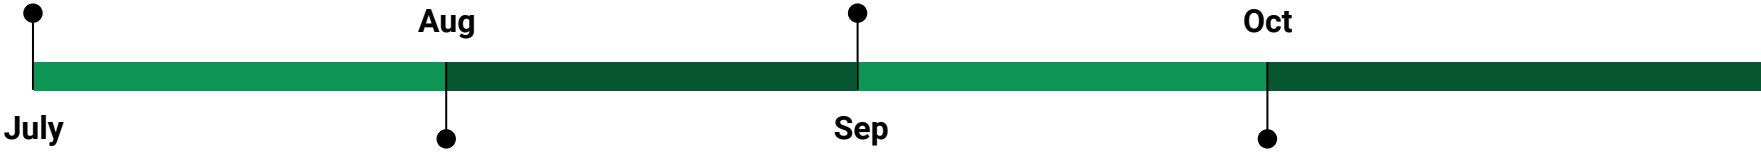

iRODS - rules, groups, vault resync
Storage Mgr - Unzip / Zip and Download drag n drop upload, show Linux homedir
Middleware - LDAP/iRODS/AFW sync
OmniSched - Permissions, enhanced connectivity

AWS - FCS
Platform - GCP
Platform - KVM
Jobs - Job Routing
UI - Job Builder

Services

- **Remote Tier 1.5 support**
  - **Hand off from HBS to Omnibond to help with:**
    - Platform Development
    - Application Integration
    - Applet AI Development
    - Researcher 1 on 1 Support
- **Project Management**
  - Omnibond can optionally provide a portion of a project manager to coordinate the overall project
- **R&D**
  - Cutting Edge Research & Development for future features

Questions?

boyd@omnibond.com

see the show