Multicloud Lustre-as-a-Service

Vinay Gaonkar, Co-Founder & Head of Products, vinay@kmesh.io

Kmesh.io

Leading the transition from centralized data lakes into distributed data ponds

Multi-Cloud

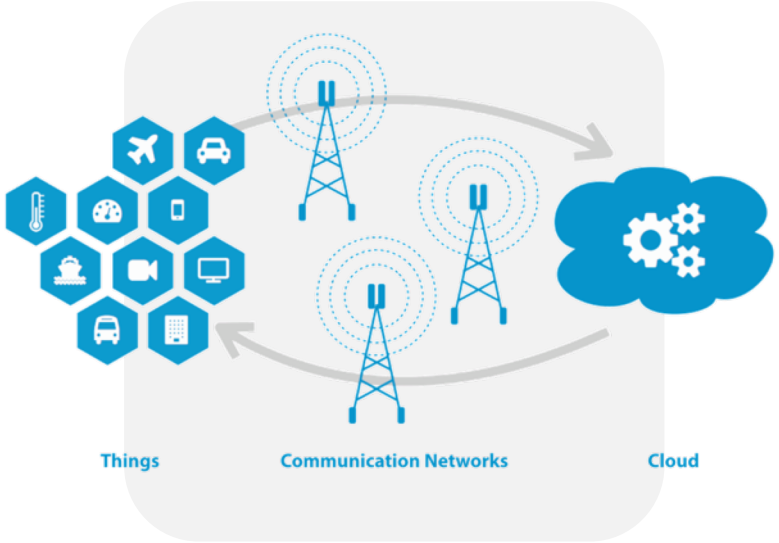

Edge

Data Sovereignty

Kmesh: True SaaS managed via Portal

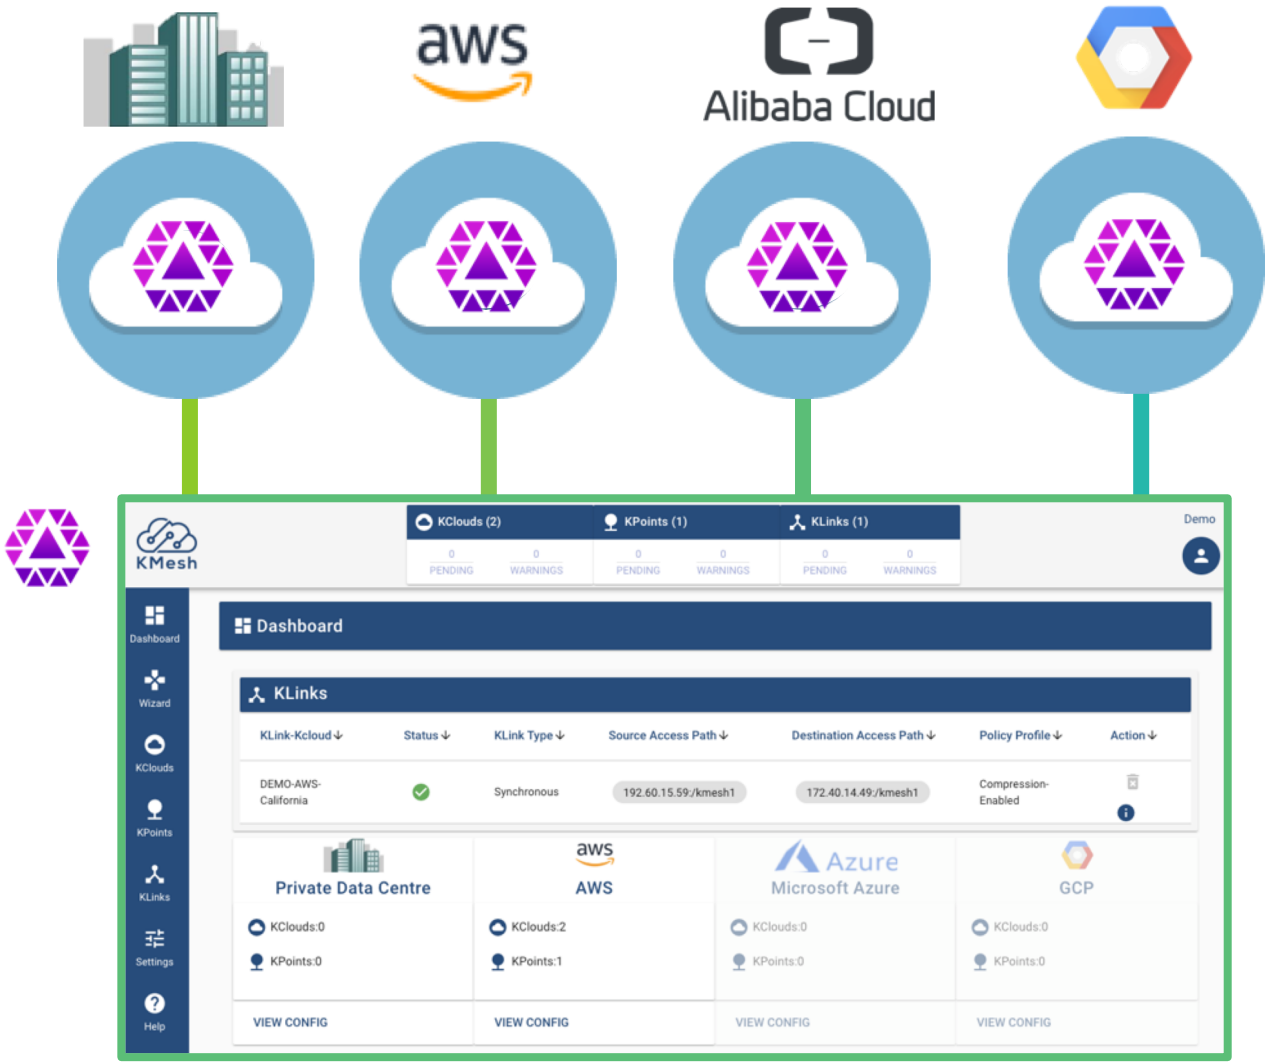

Real-time Data Mobility

Onprem, Cloud, Edge

Policy driven SaaS

Kubernetes Data Orchestration

Lustre-powered for HPC, AI, ML

The Leading Multi-Cloud Lustre-as-a-service Filesystem for HPC,AI, ML

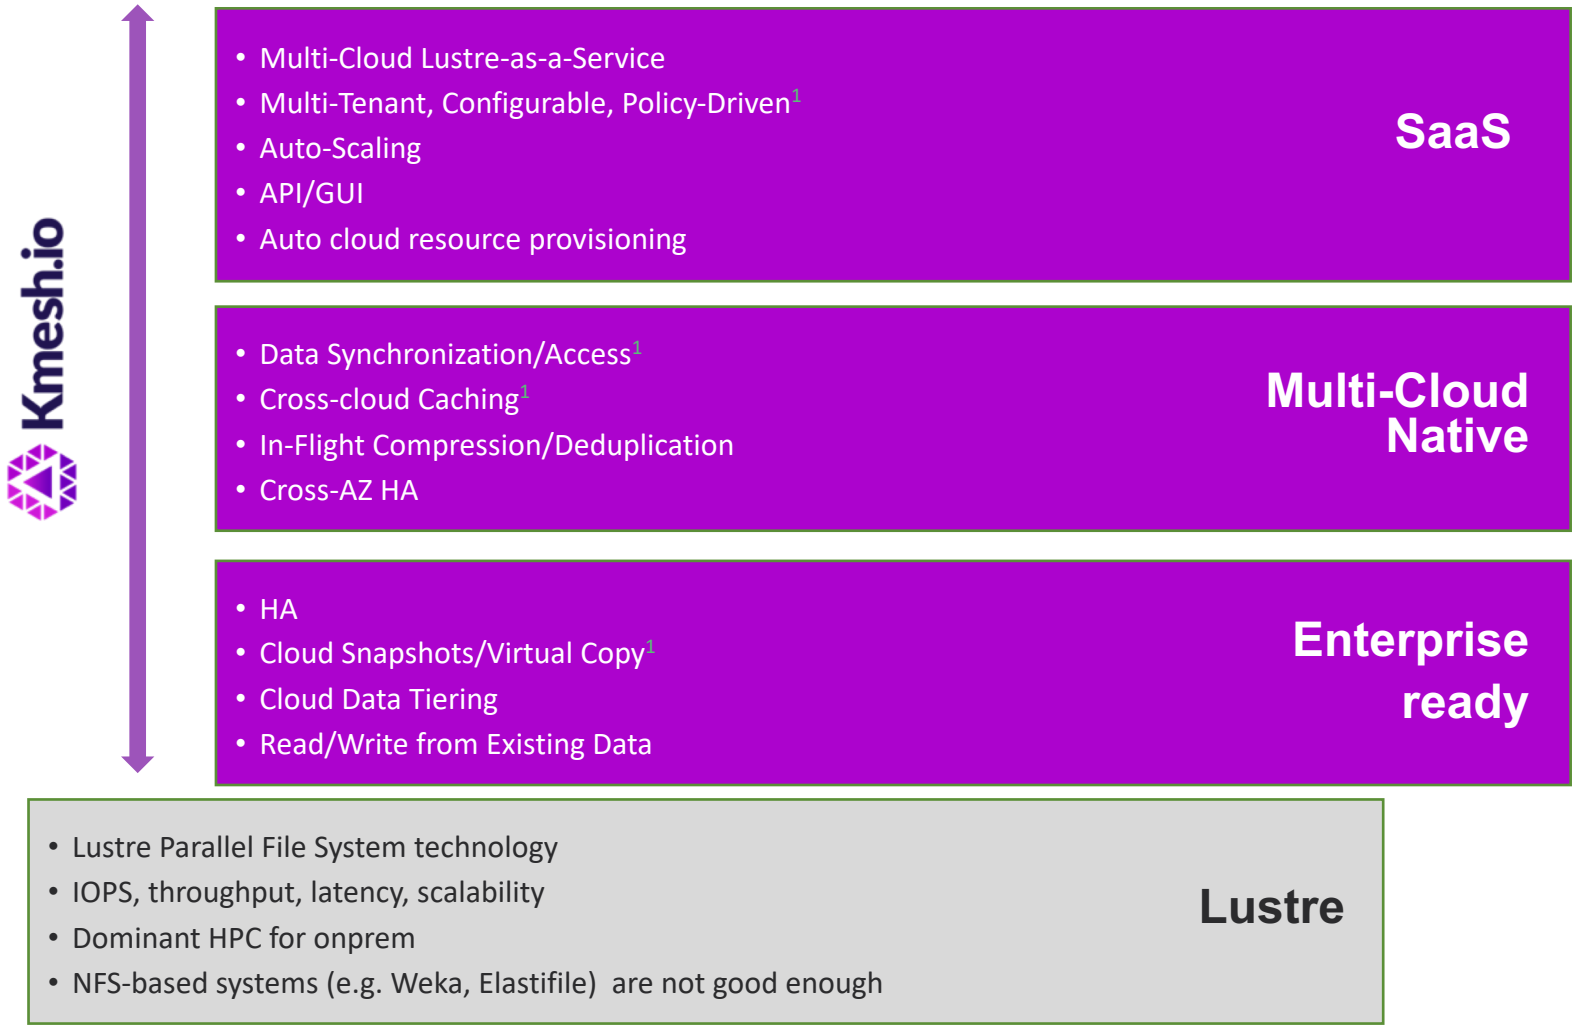

Kmesh: True SaaS managed via Portal

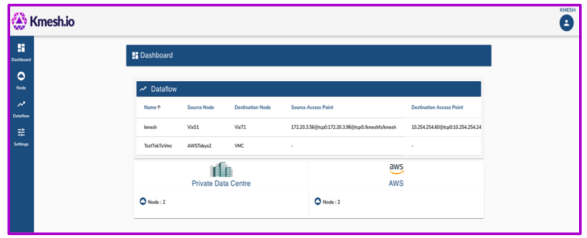

Kmesh.io

SaaS

1 User creates nodes, data flows, policies

Corp Data Center

Public Clouds

aws
Azure
Google Cloud Platform

APP

Run the app on Kmesh filesystem 3

/mydata

File Storage

Ingest/write data

2 Kmesh SaaS deploys the agents

Kmesh Data Flow

WAN/ SDWAN
Encryption, Compression

Kmesh Data Node

Kmesh data node is deployed as a VM
vmware

APP

3 Run the app on Kmesh filesystem

/mydata

File Storage

Kmesh ingests/write data

Object Storage

Kmesh Data Node

Kmesh data node is deployed on most-optimal instances and storage

Kmesh Dataplane Architecture

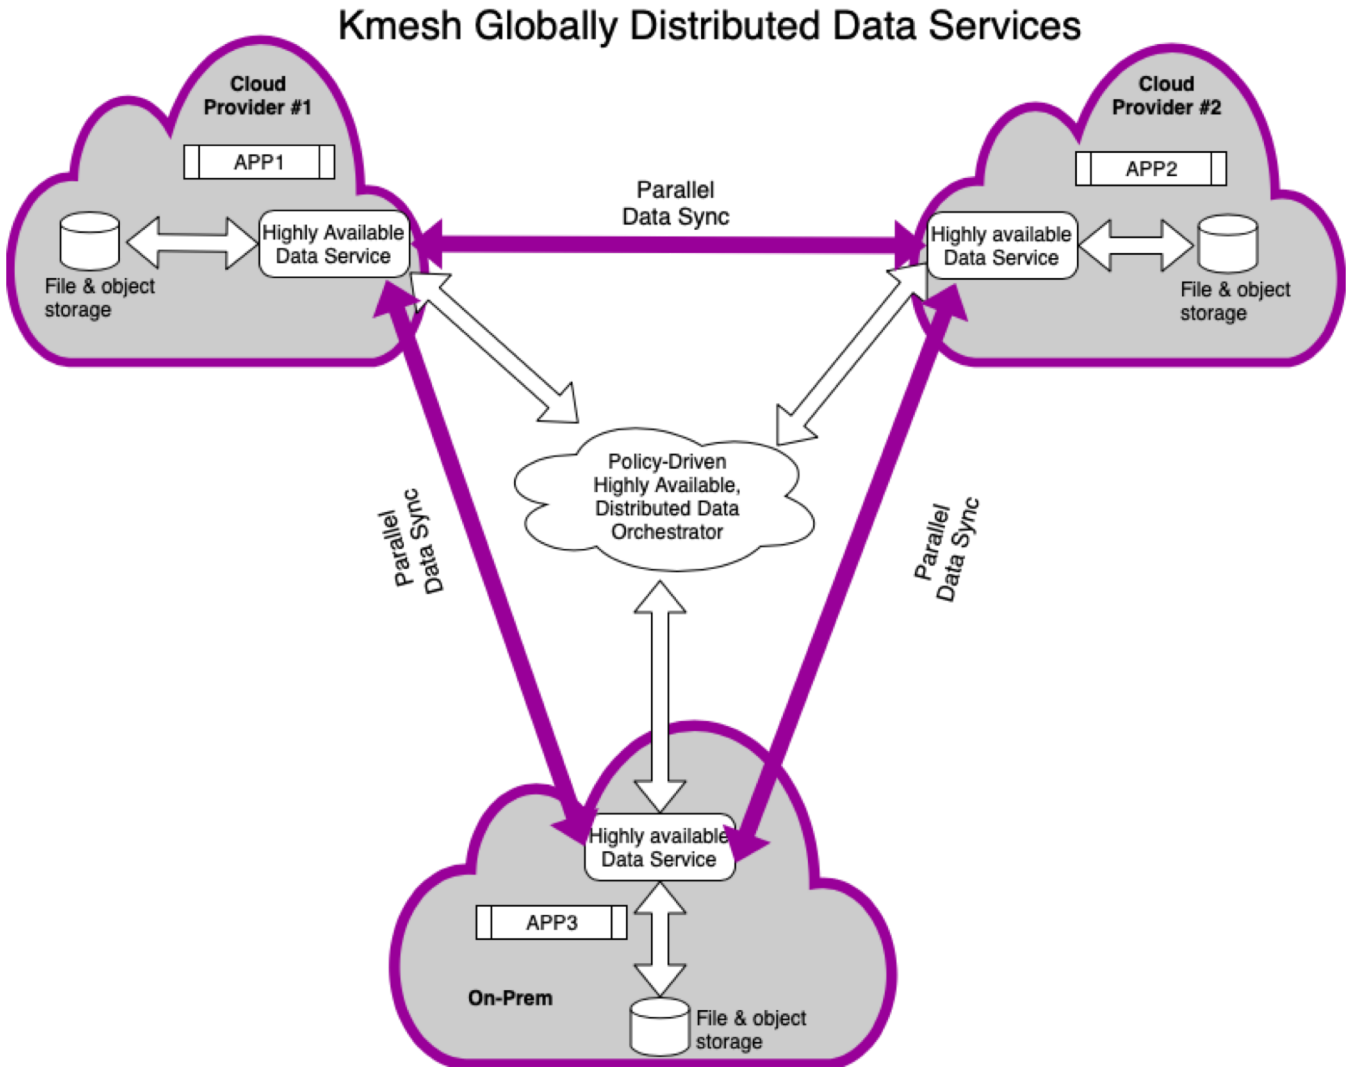

Kmesh Data Node

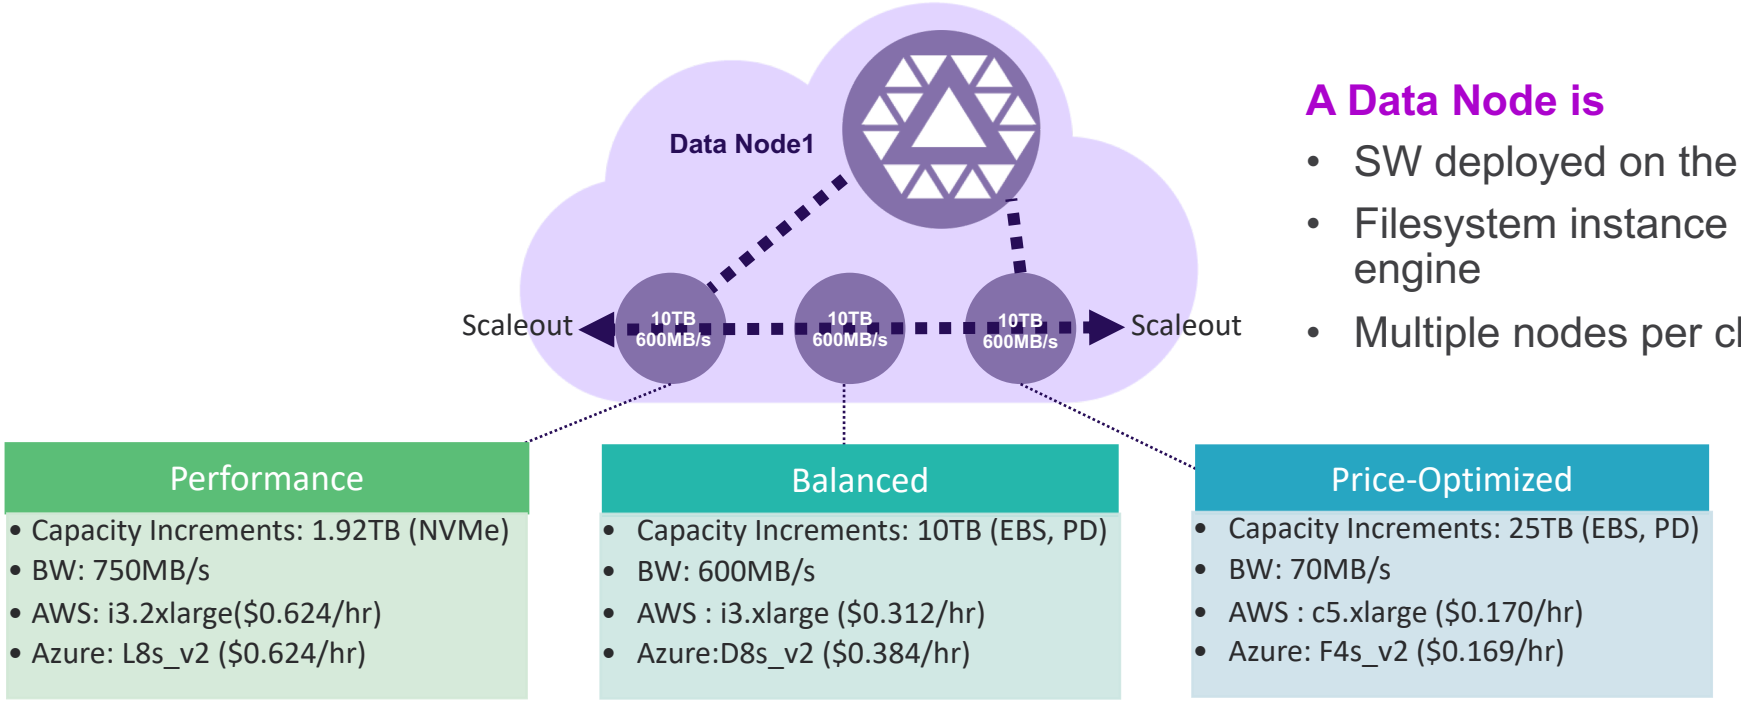

* Performance and instance/VM costs on AWS and Azure are subject to change.

A Data Node is

- SW deployed on the cloud
- Filesystem instance & data synchronization engine
- Multiple nodes per cloud

Cost-optimized provisioning

- Different Node Types to meet app requirements
- Auto-provisioning & auto-scaling of nodes
- Data Tiering

Features

- At-rest compression & deduplication
- HA and Non-HA
- Snapshots/Virtual Copies

Kmesh Multi-Cloud Snapshots

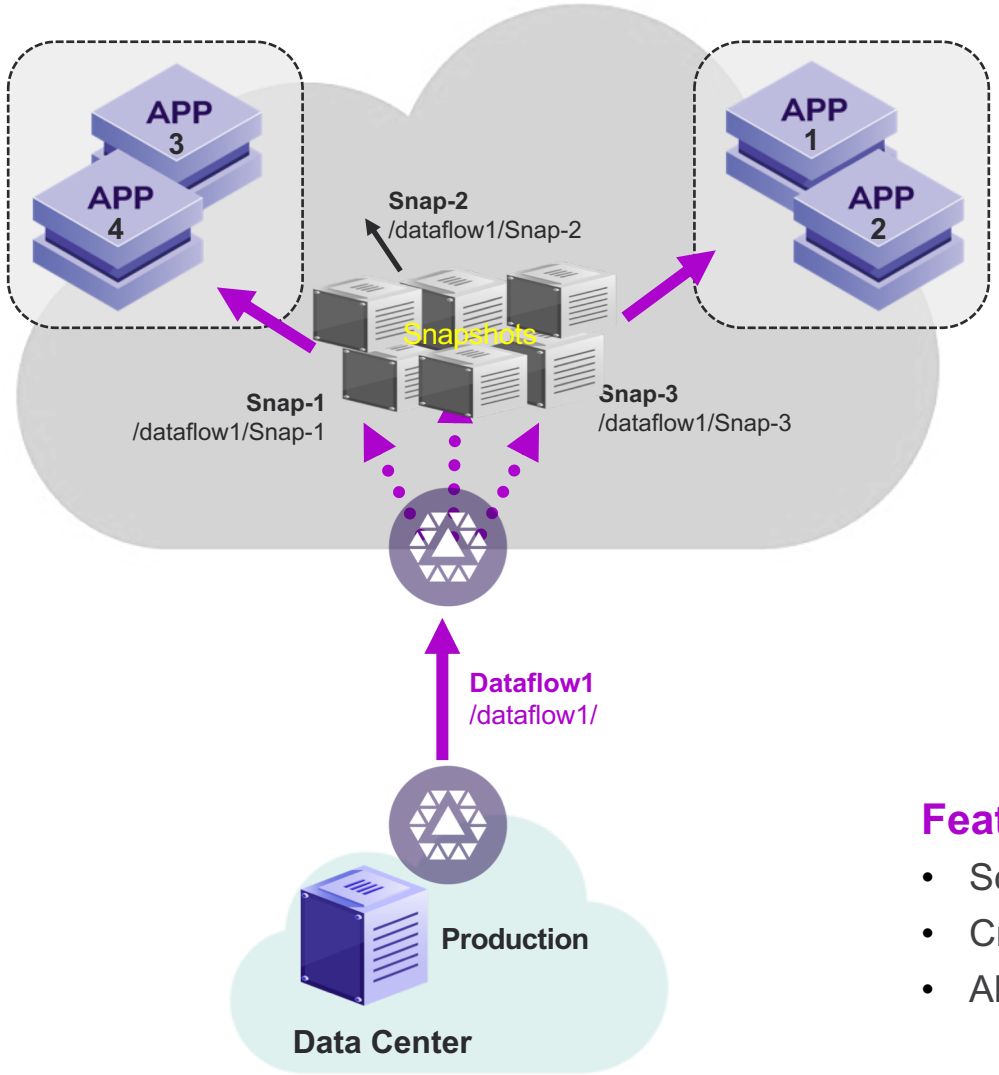

Snapshots

- Single & Multi-cloud Snapshots
- Cross-cloud App consistency for variety of apps
- User/group/app based access to snapshots

Features

- Scheduled (Daily/weekly.....)
- Cross-cloud storing of snapshot
- Ability to mount remotely via dataflow

Kmesh Multi-Cloud Virtual Copies

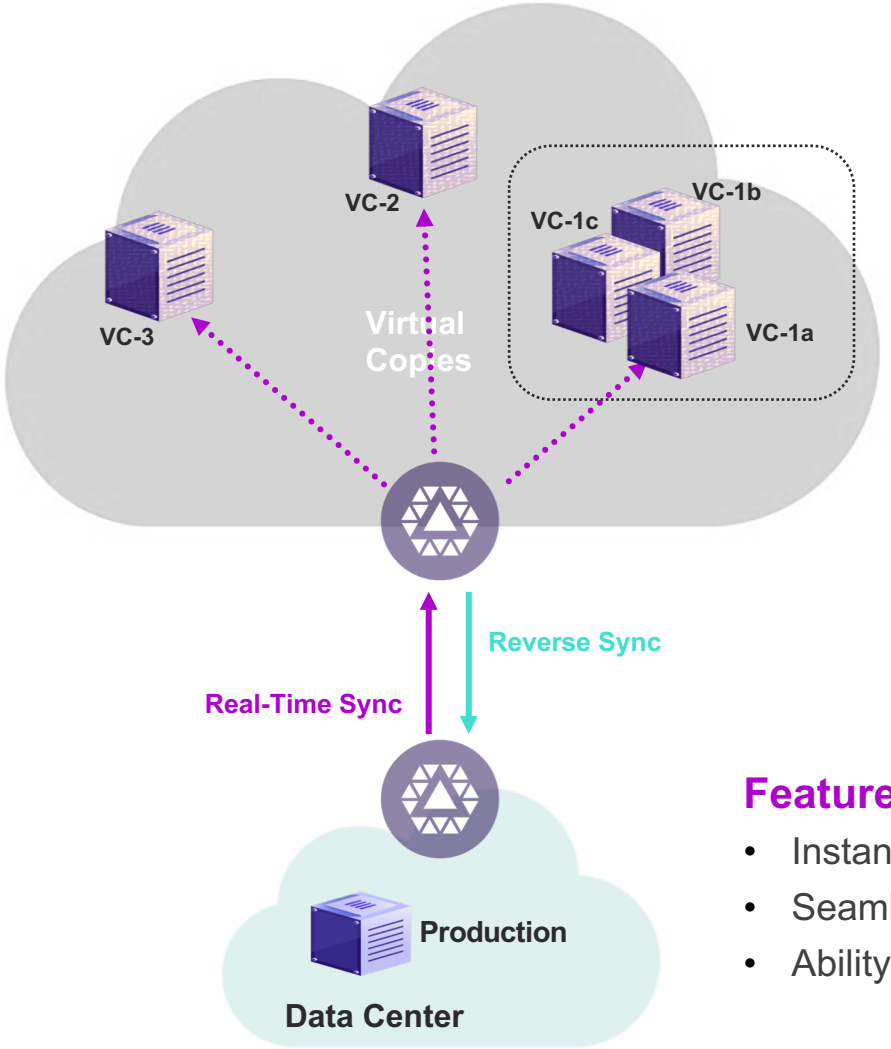

Read-write Virtual Copies

- Space efficient & instantaneous
- App-consistent copies
- App/user/group level access management

Features

- Instantaneous refresh to production
- Seamless & efficient reverse-sync
- Ability to mount remotely via dataflow

Data Tiering

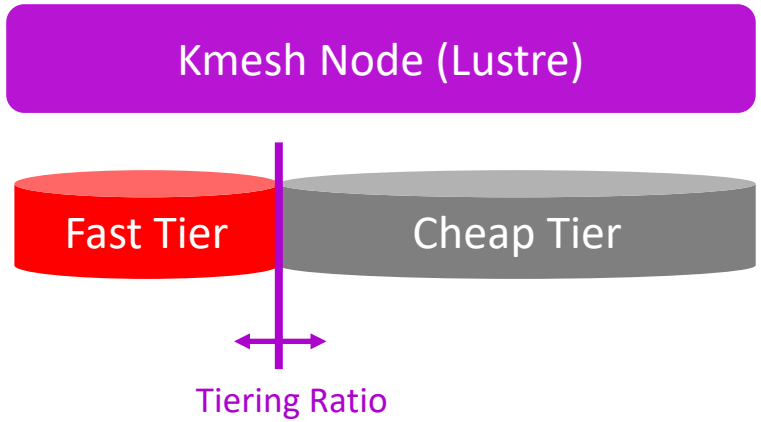

- Tiering ratio set during node creation
- When fast tier fills. Least accessed files are moved to cheap tier
- Files are brought into fast tier if not present

Kmesh Data Flows

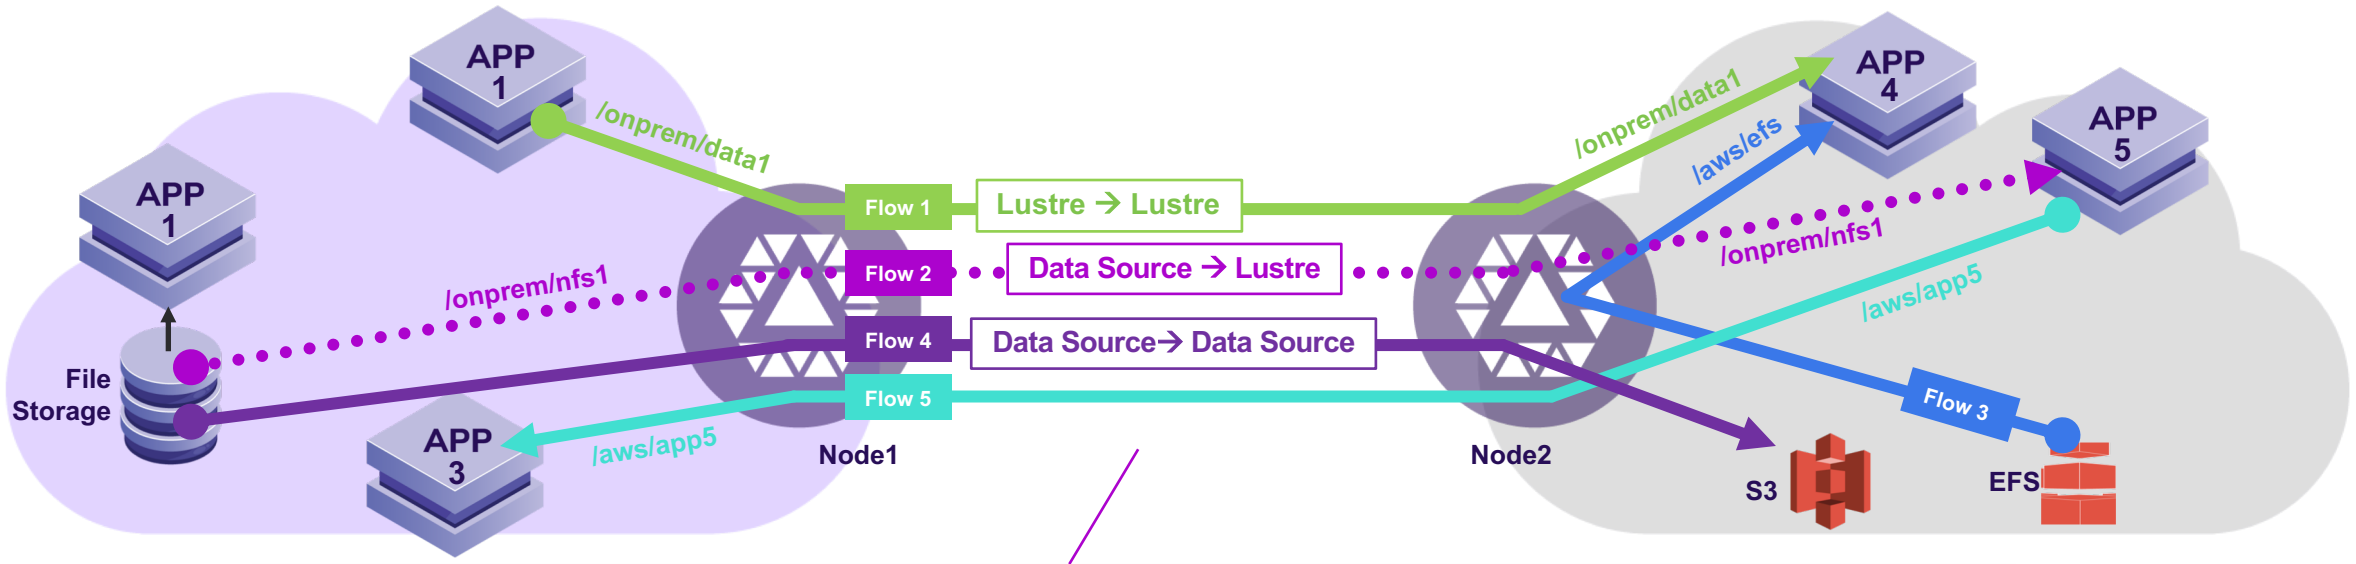

Data Flow is a data connection accessed as a mount point

Filesystem ↔ Filesystem Data Source ↔ Filesystem Data Source ↔ Data Source

- **Real-Time Data synchronization**
  - Sync, Async and Multi-master
- **Remote Data Access**
  - Intelligent cross-cloud caching
- **Scalable**
  - Efficient, parallel stream and flexible

- In-flight compression & encryption
- Temporal data access for sensitive data
- *Data protection with snapshots*
- *Virtual copies (cloud clones)*

Different HPC Cloud Bursting Models

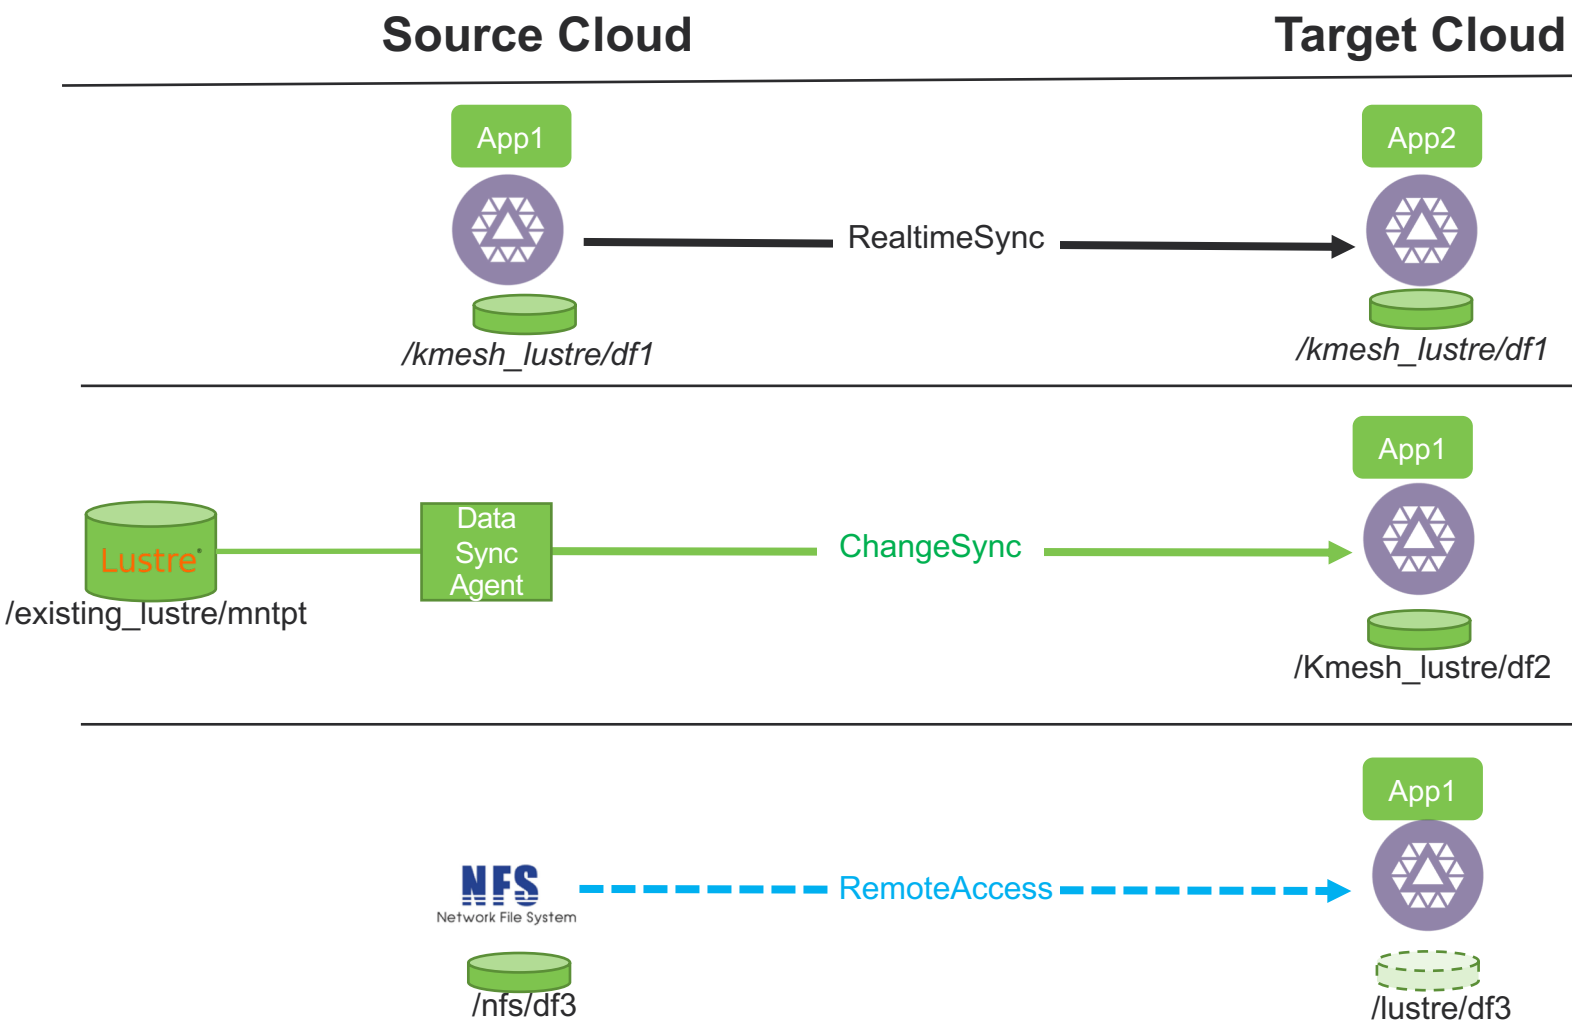

Hybrid Cloud Kmesh

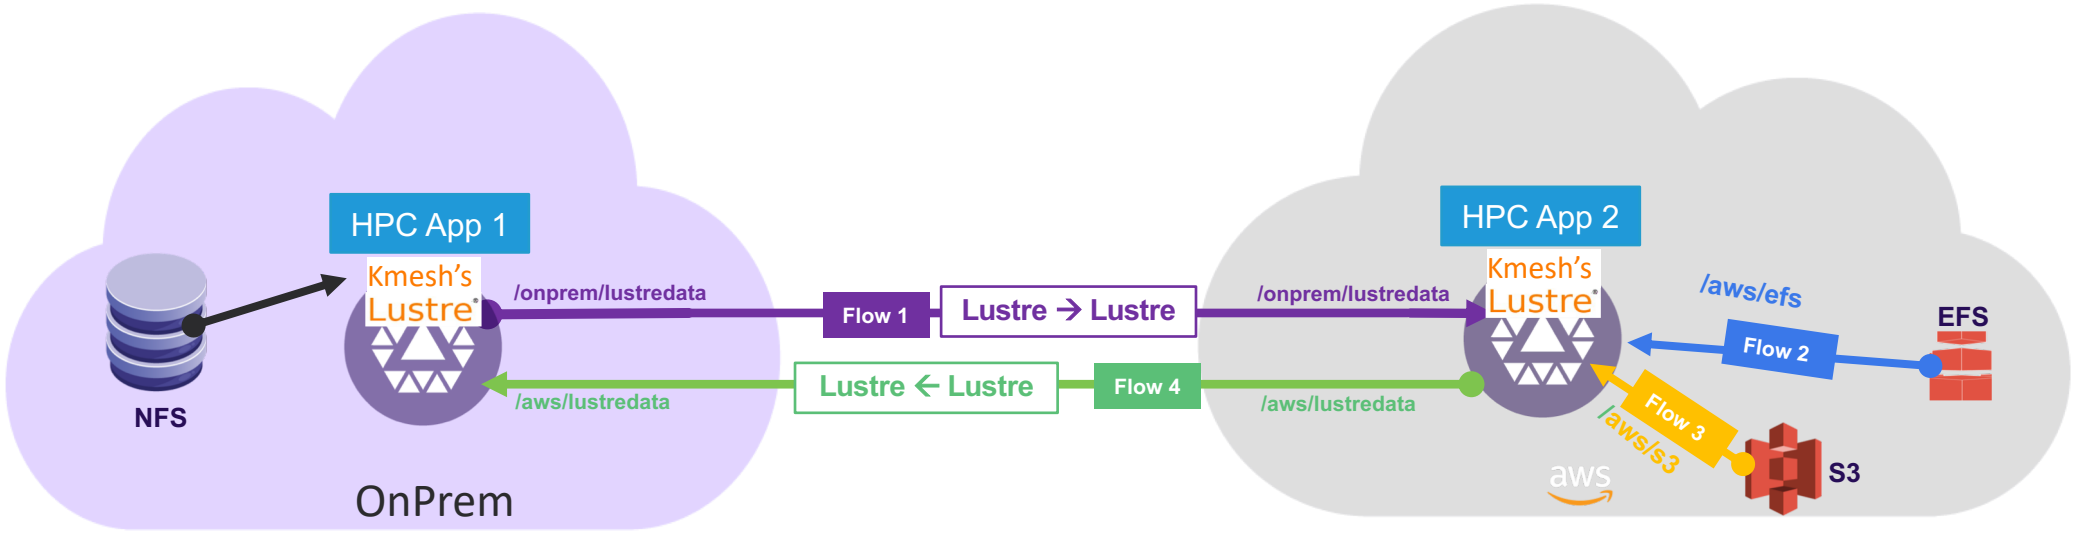

Ingest from existing Lustre

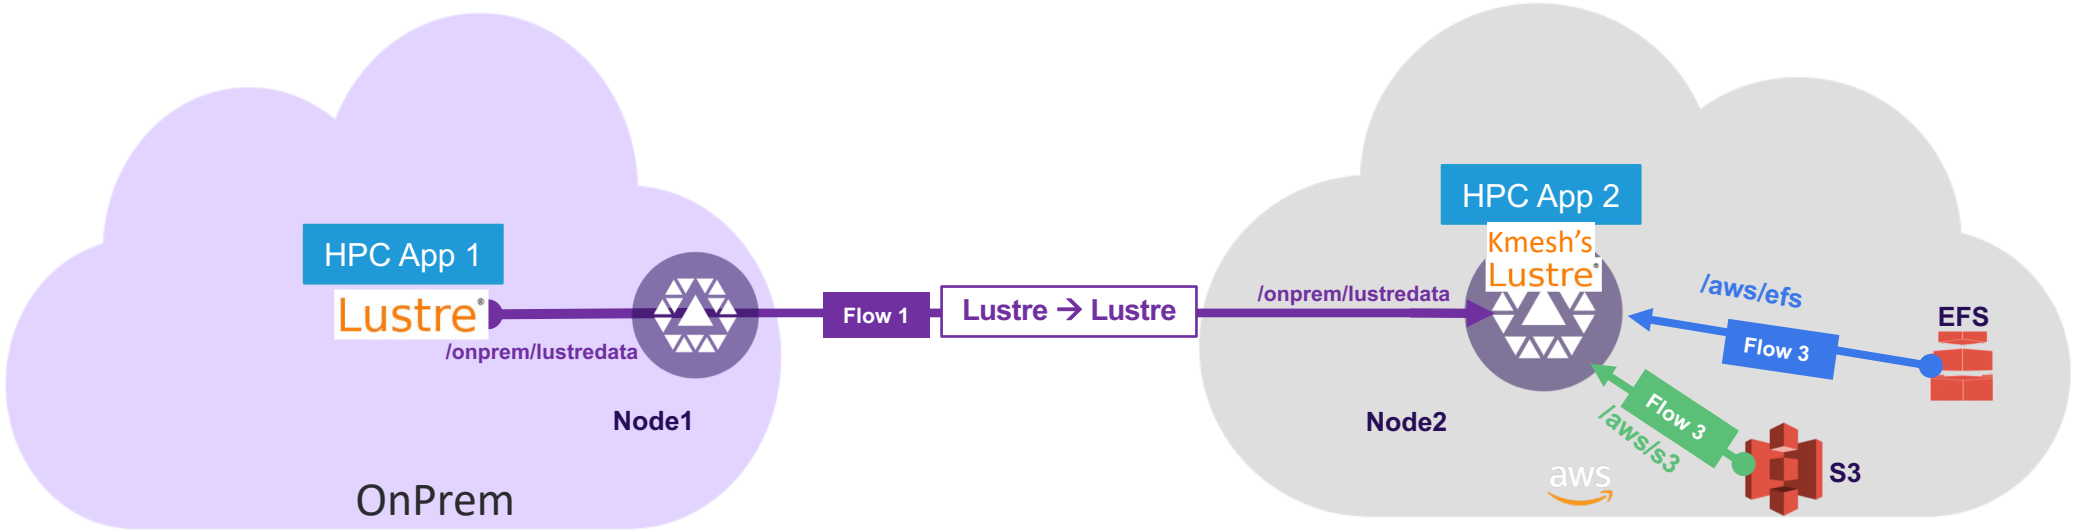

Kmesh Lustre on Cloud remotely accessing data

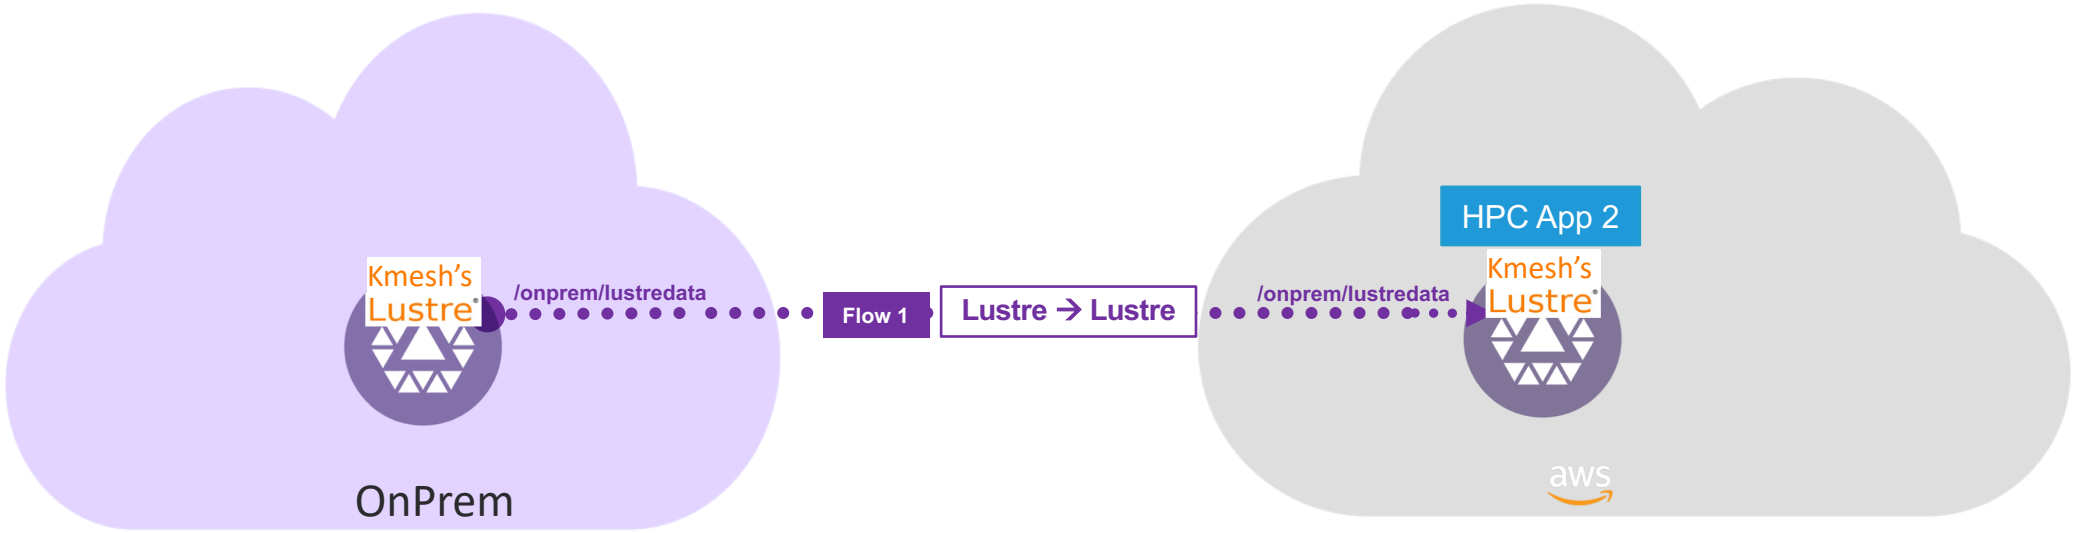

FREE POCs for the LUG community!
Stop by our table OR vinay@kmesh.io

We are hiring Lustre Engineers!

THANK YOU.

Kmesh.io