

The HDF Group Google Cloud HPC Blueprint

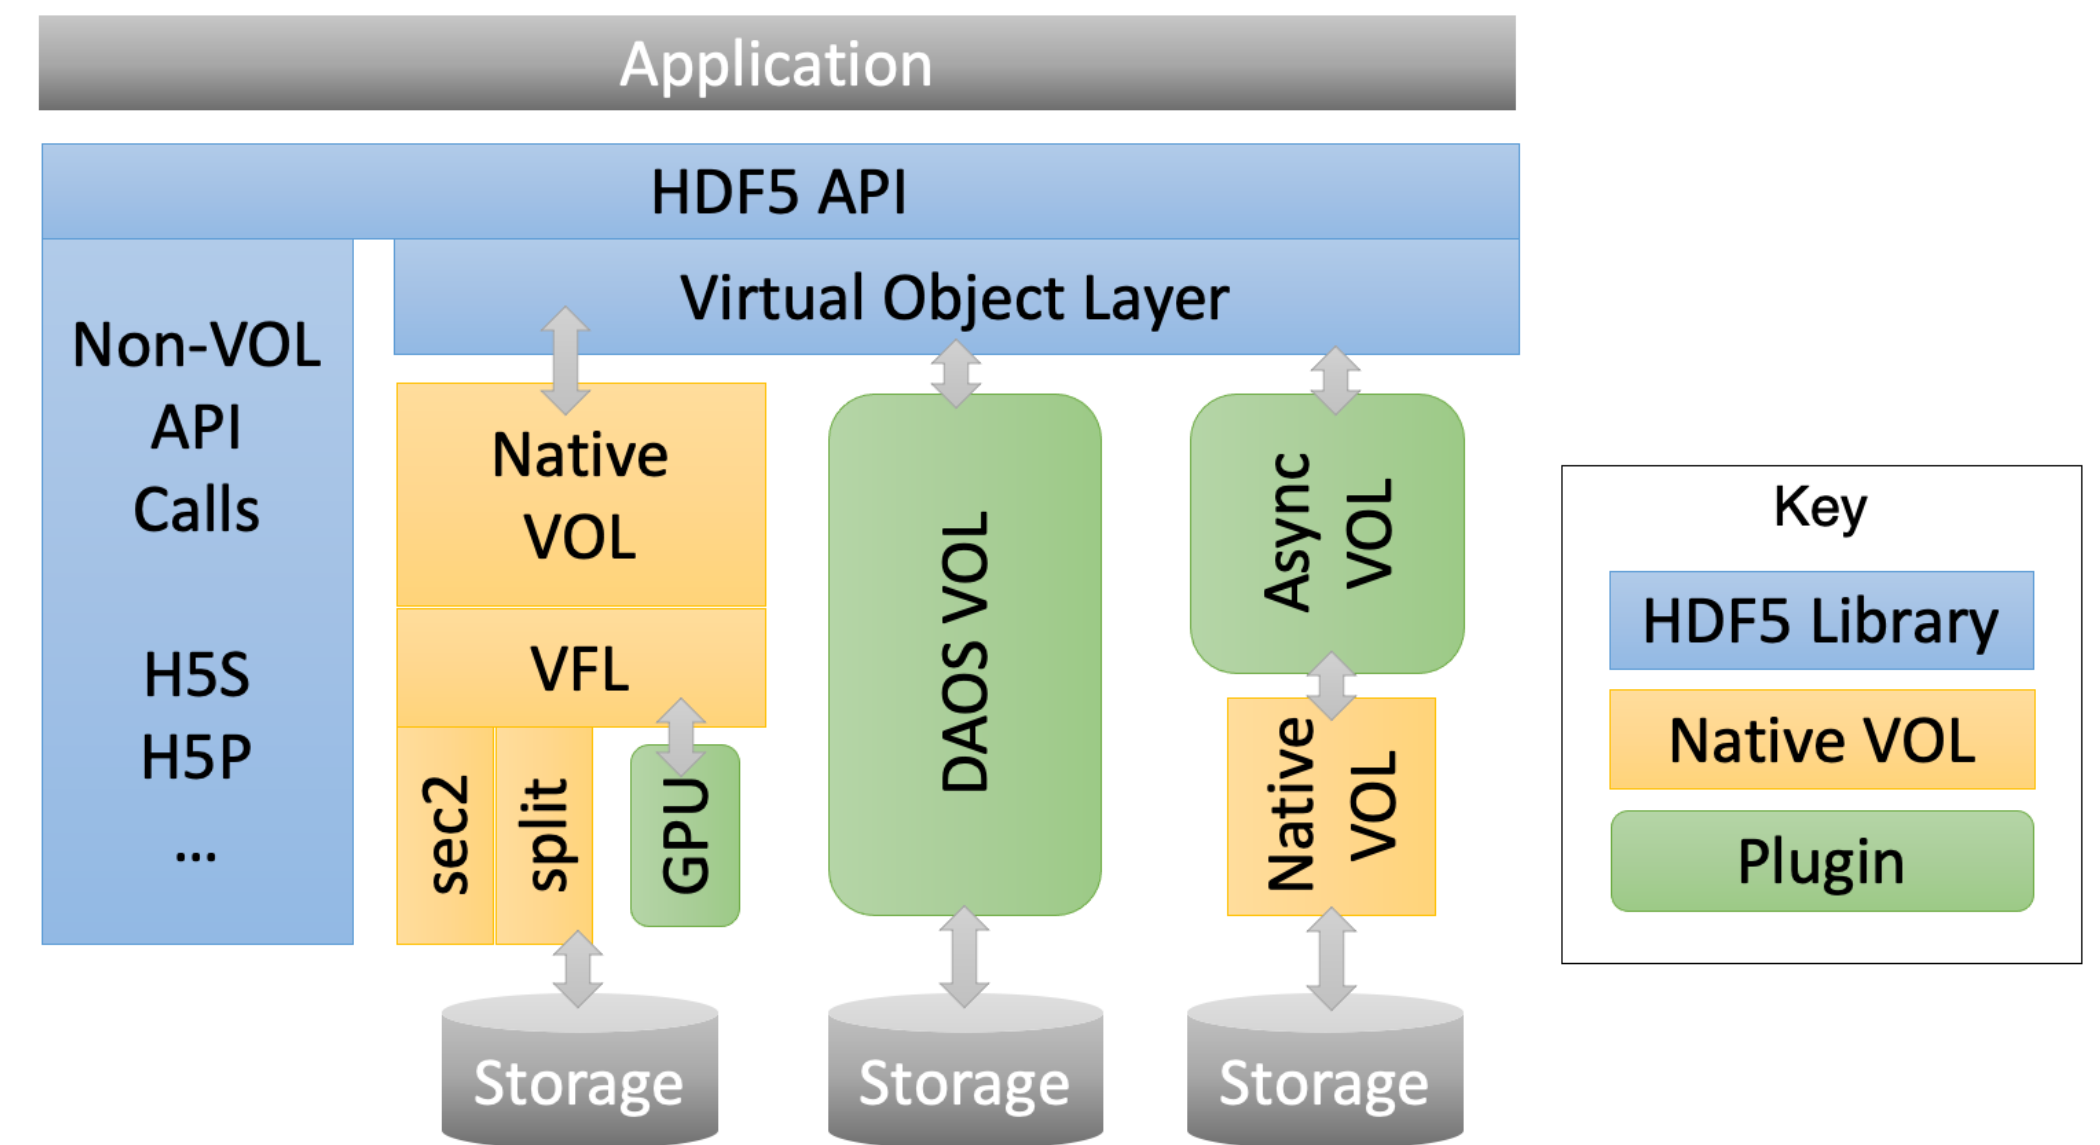
Aug 2023



Glenn Song
The HDF Group

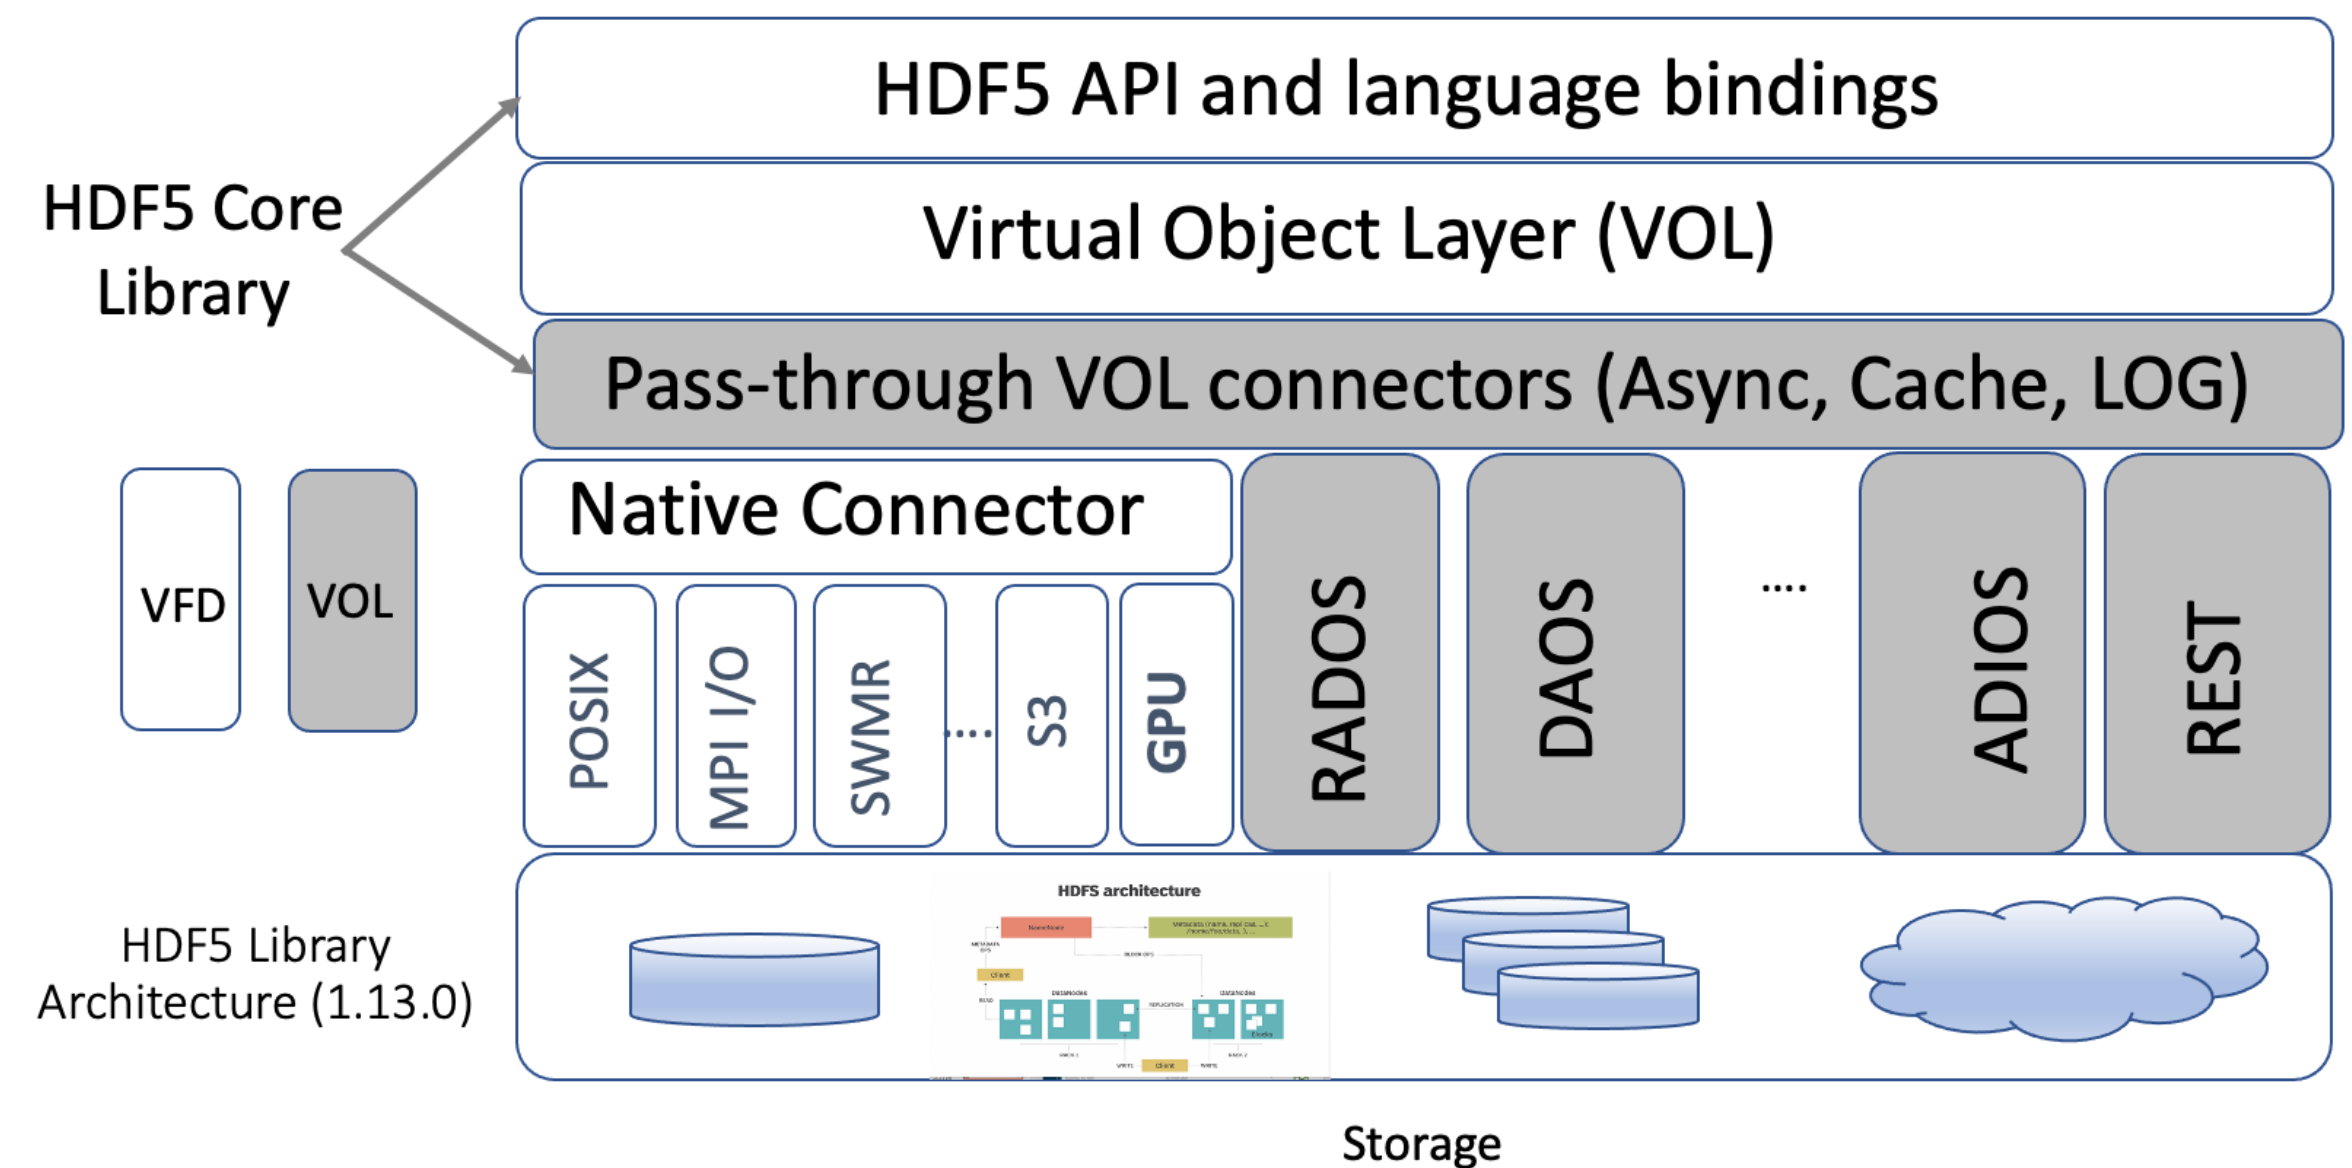
VOL (Virtual Object Layer)

- The VOL exists just below the public API
- It intercepts all **storage** calls that might access objects in an HDF5 container and forwards them to a VOL connector
- Once in control of a VOL connector, the connector does everything it needs to do, then passes control back to the library



VOL Connectors

- Passthrough and terminal
 - Passthrough: can invoke other VOL connectors underneath itself
 - Terminal: don't pass to other VOL connectors, they are designed to map HDF5 file objects/metadata to storage
- Users can define their own VOL connectors and test using our VOL/API test suite



DAOS VOL Connector



- VOL connector for HDF5 that allows for direct interfacing with the DAOS (Distributed Asynchronous Object Storage) system
- Bypasses both MPI I/O and POSIX for efficient and scalable I/O
- Allows for new features like creating objects in parallel, key-value store objects, data recovery, asynchronous I/O, etc.
- If you are building from source code, must use Cmake
- <https://github.com/HDFGroup/vol-daos>
- However, we will be building this using Spack

HPC in Google Cloud



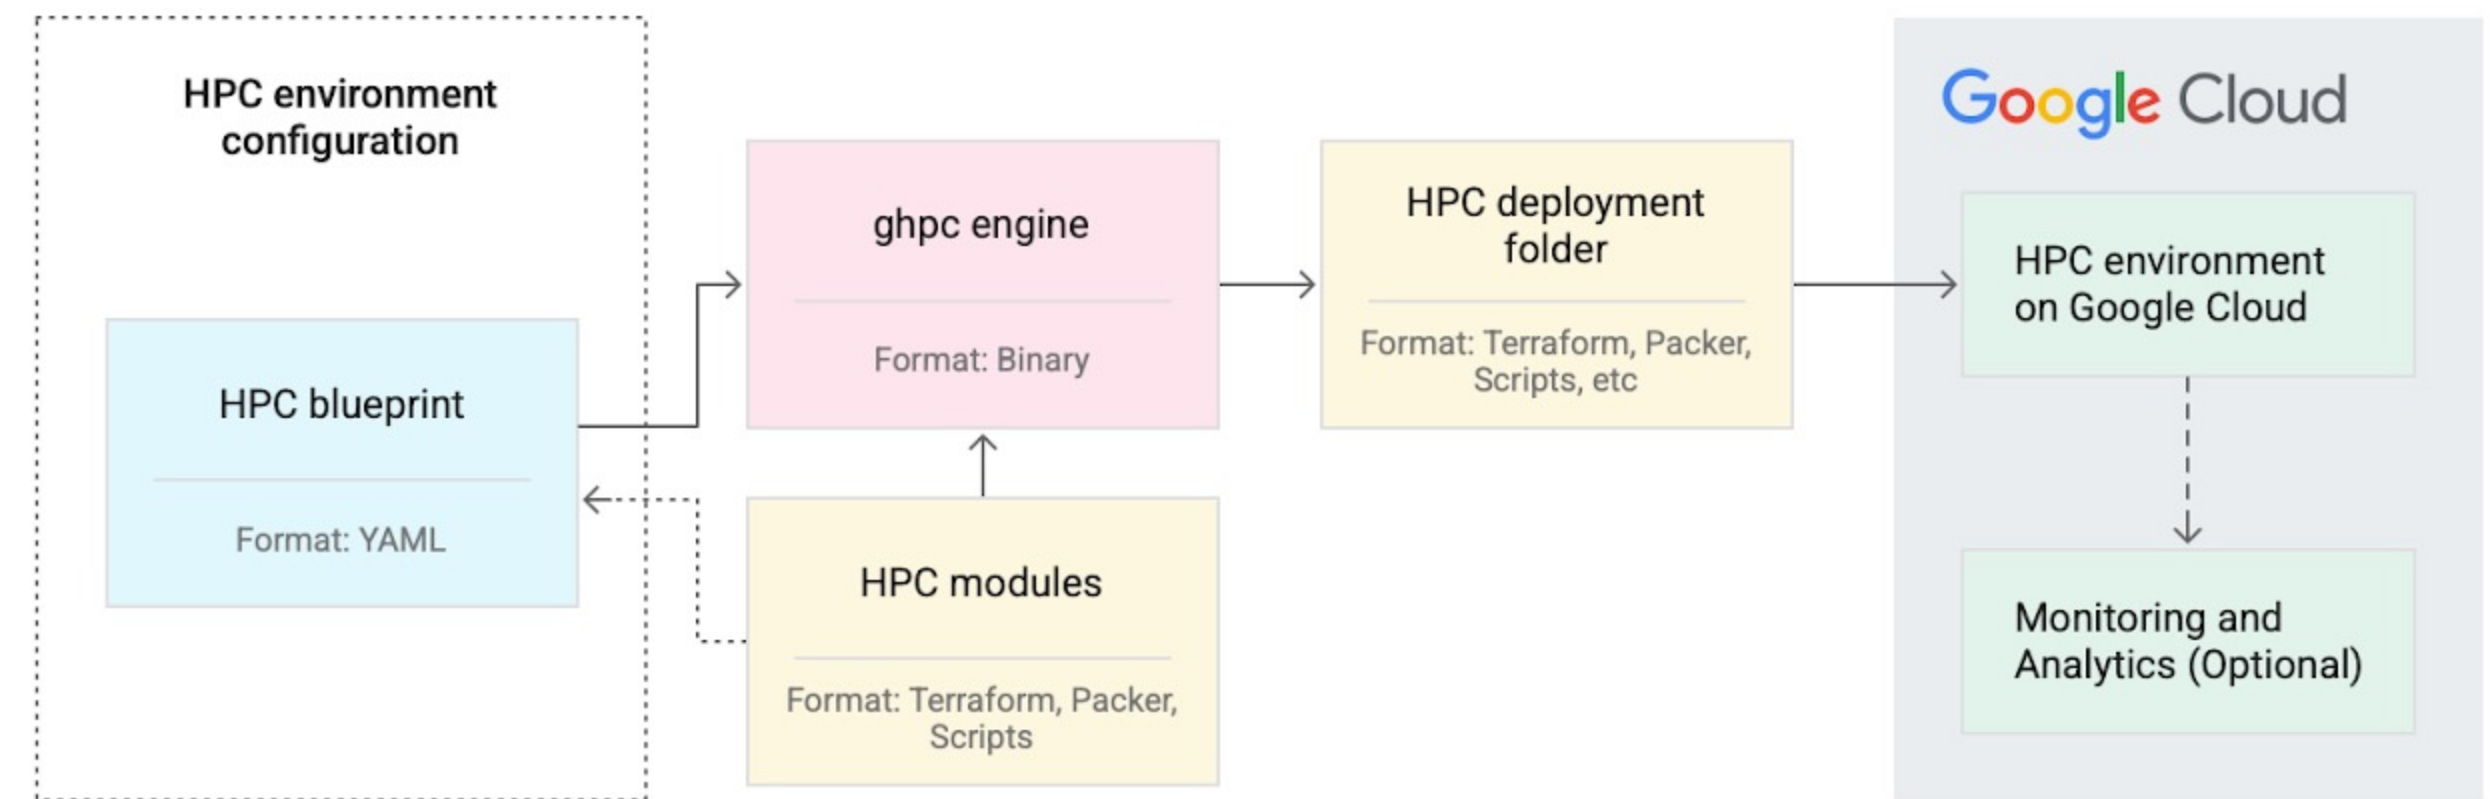
- As cloud-based services become more and more widespread and available, HPC services are also becoming available in the cloud
- Google Cloud's HPC Toolkit is meant to make this transition easier
- HPC blueprints are reusable configurations created using a YAML file that describes the exact HPC environment you want to deploy
 - Uses just the HPC Toolkit
- Allows users to easily deploy and destroy clusters quickly



Google Cloud

Blueprints

- Has three parts: HPC blueprint name, deployment variables, deployment groups
 - Blueprint name is important for tracking usage and costs
 - Deployment variables are used to set values specific to a deployment
 - Deployment groups define a group of modules that will be deployed together
- Can use Spack, Slurm, etc. in modules



Our Blueprint



- Created in collaboration with Google
- Will set up HDF5, CMake
- Can set up a DAOS client and server, and can be extended to an entire cluster
- Will eventually set up the DAOS VOL Connector using Spack

Common Issues and Concerns

- Quota issues
 - Biggest issue when working with Google Cloud HPC
 - Hard to tell when running into quota issues sometimes
- When a blueprint hasn't been deployed properly, it can be hard to shut everything down properly
 - Will need to go look at all the resources that are still running
- Should always check resources page to make sure everything is destroyed if necessary
 - Can result in unnecessary costs otherwise

THANK YOU!

Questions & Comments?