



Cloud Native HPC

Tanmoy Palit, *Technical Lead*

Johan Seland, *PhD, Team Lead Cloud Services*

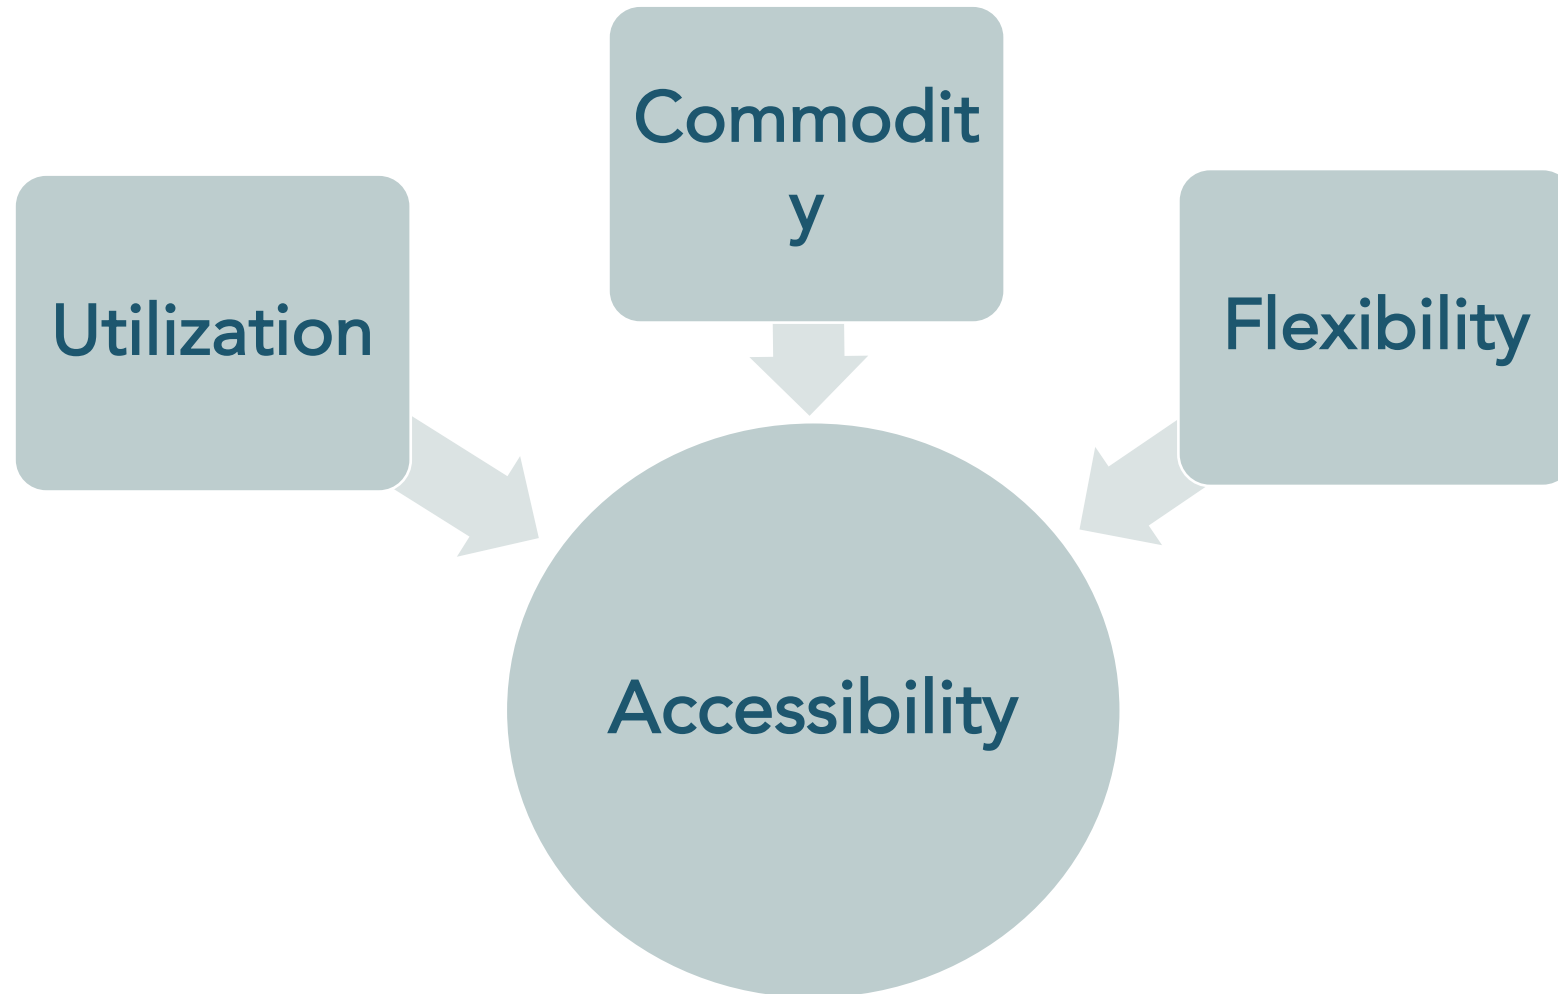
Alberto Melacini, *Senior Cloud Architect*

Gudbrand Tandberg, *Software Developer*

Kostya Deev, *Senior Cloud Developer*



HPC in the Cloud



How Do We Get There?



INTERACTIVE

BATCH

High Fault Tolerance

On Demand Scaling

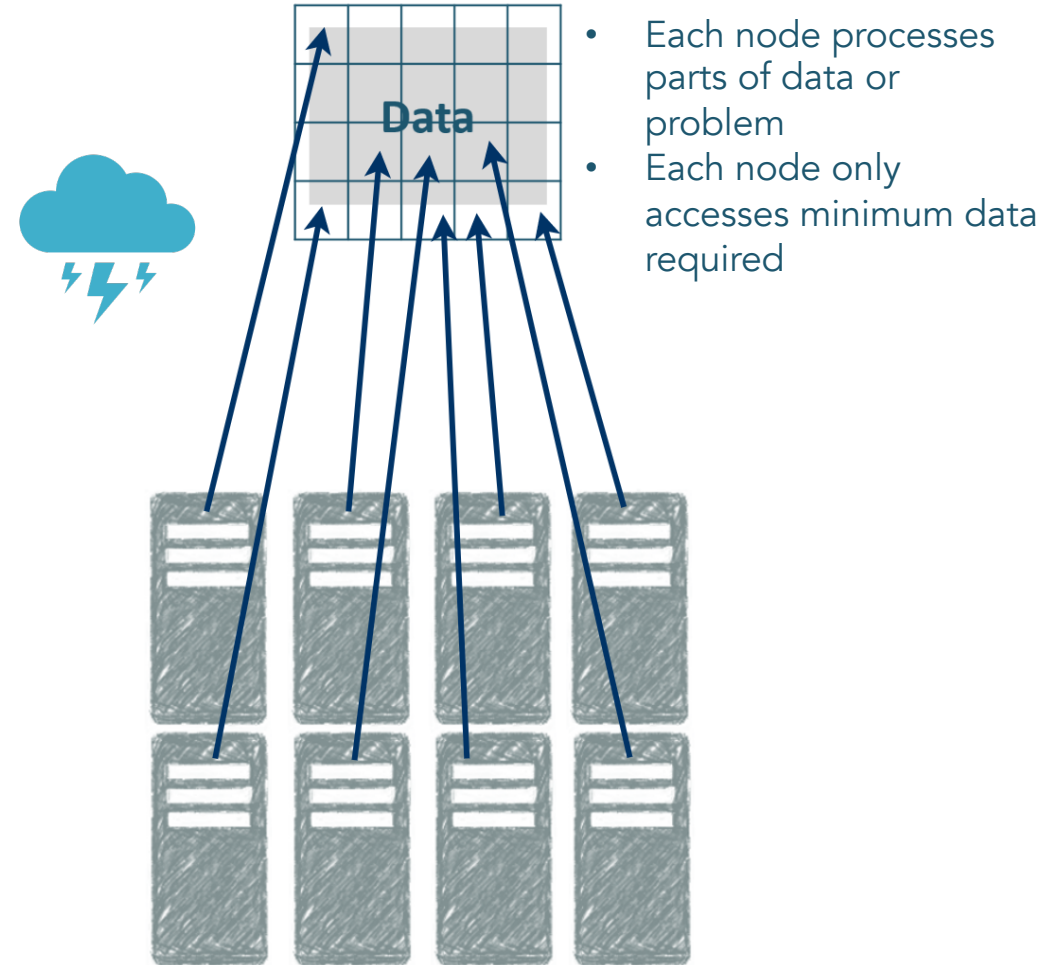

Distributed Workflows Instead Data

Efficient Compute Engine

Utilize Compression

Scalable, High Performance Data
Store

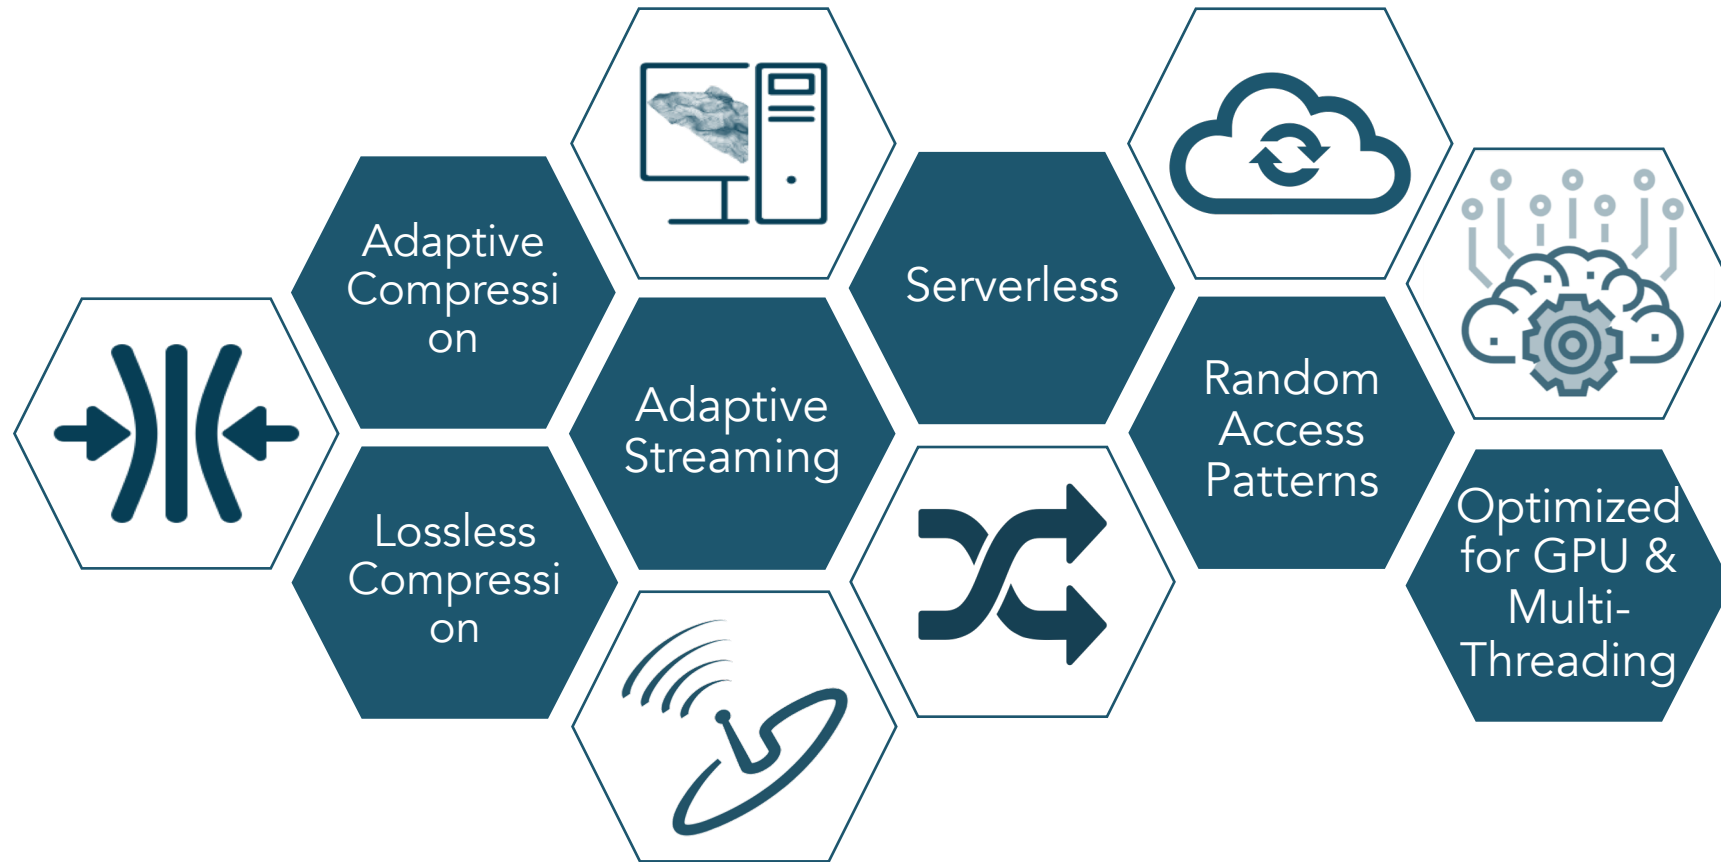
Cloud Native Scalability



Volume Data Store (VDS)



Compress raw and interpreted volumes of seismic datasets that are adaptable and scalable depending on business needs, workflows, and visualization requirements:




Evolution of VDS API

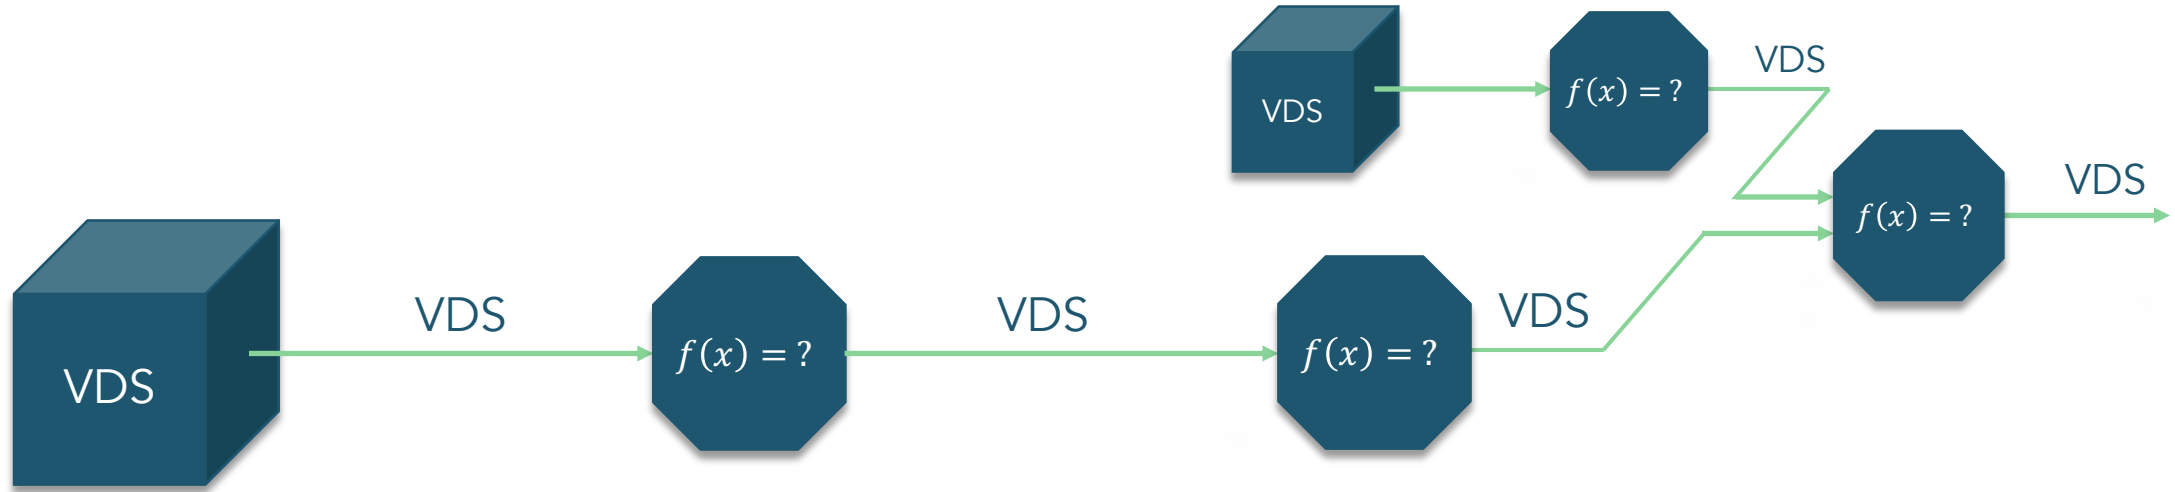


VDS >>
VDSREMOTE >>
WORKFLOWREMOTE

- Cloud Ready
- Native Object Store Format
- REST API
- Compressed Data Transfer
- Dynamic Compute Workflows



Bluware COMPUTE



VDS delivers	Compute Plug-in	Chained Compute	Combined Compute
<ul style="list-style-type: none">▪ Fast▪ Random Access▪ Adaptive Streaming	<ul style="list-style-type: none">▪ VDS Input(s)▪ Performs Calculation▪ Outputs VDS	<ul style="list-style-type: none">▪ Acyclic Graph of Actions▪ Calculate Chain on Execution▪ High Reuse	<ul style="list-style-type: none">▪ Powerful Interactive Workflows▪ Complex Problem Solving▪ Dynamic Orchestration



Distributed Compute Framework



Compute Server



Load Balancer



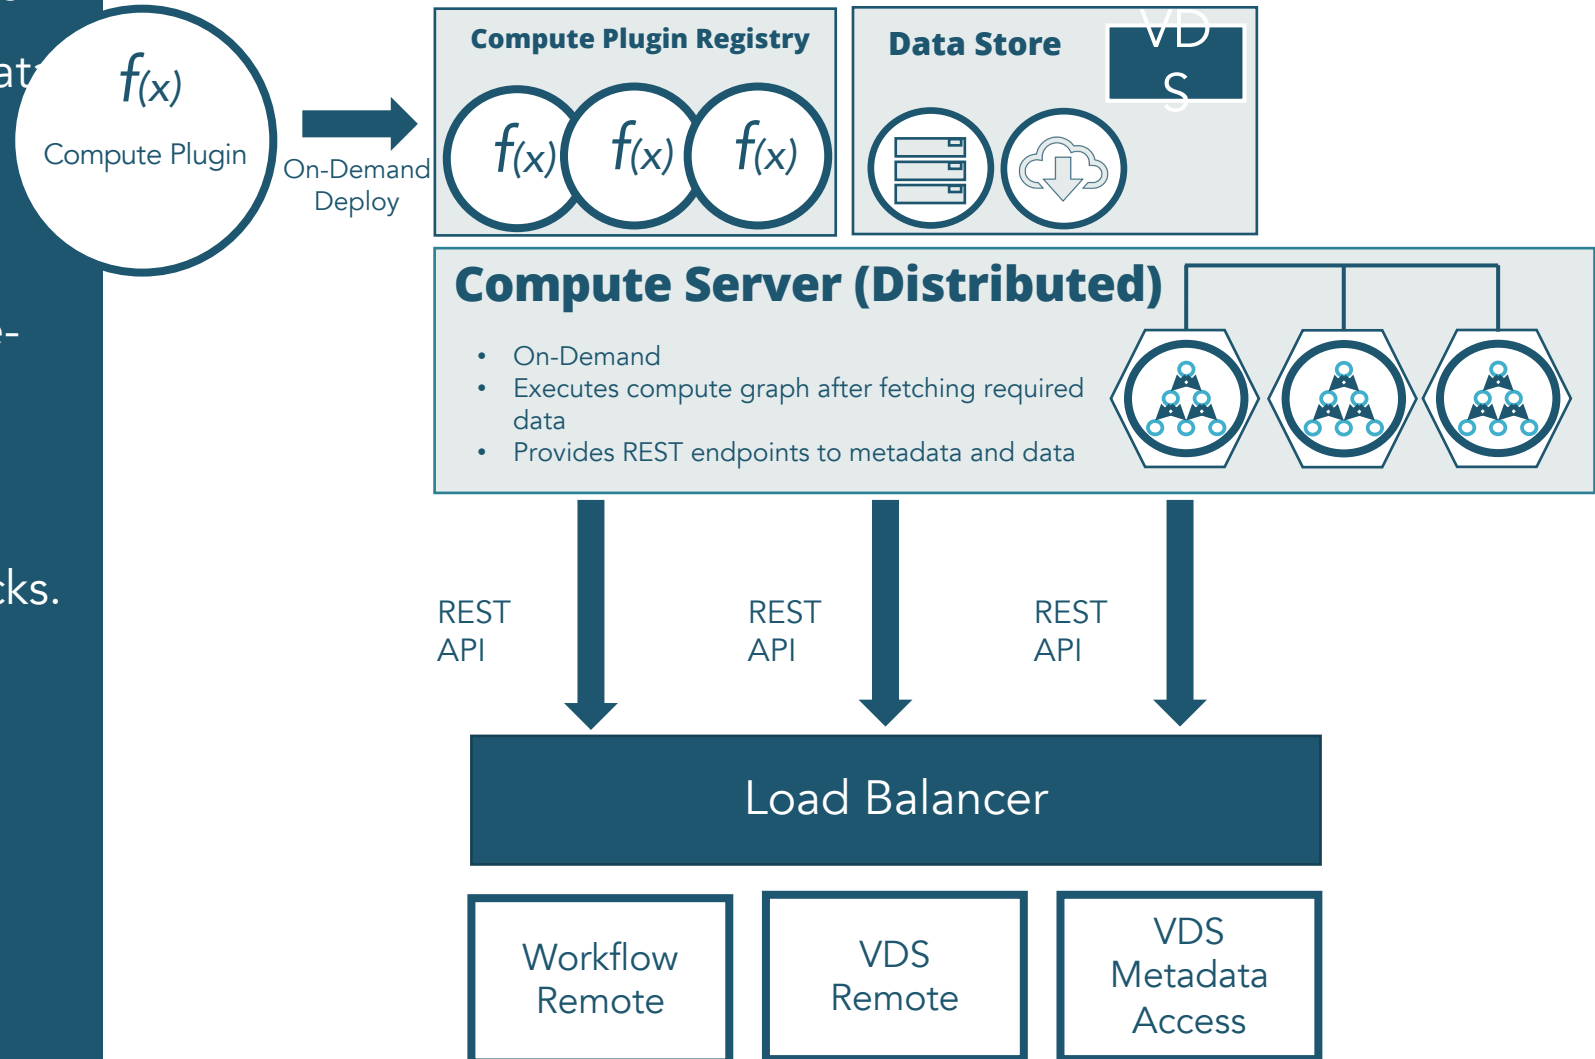
Workflow
State Manager



Session
Manager

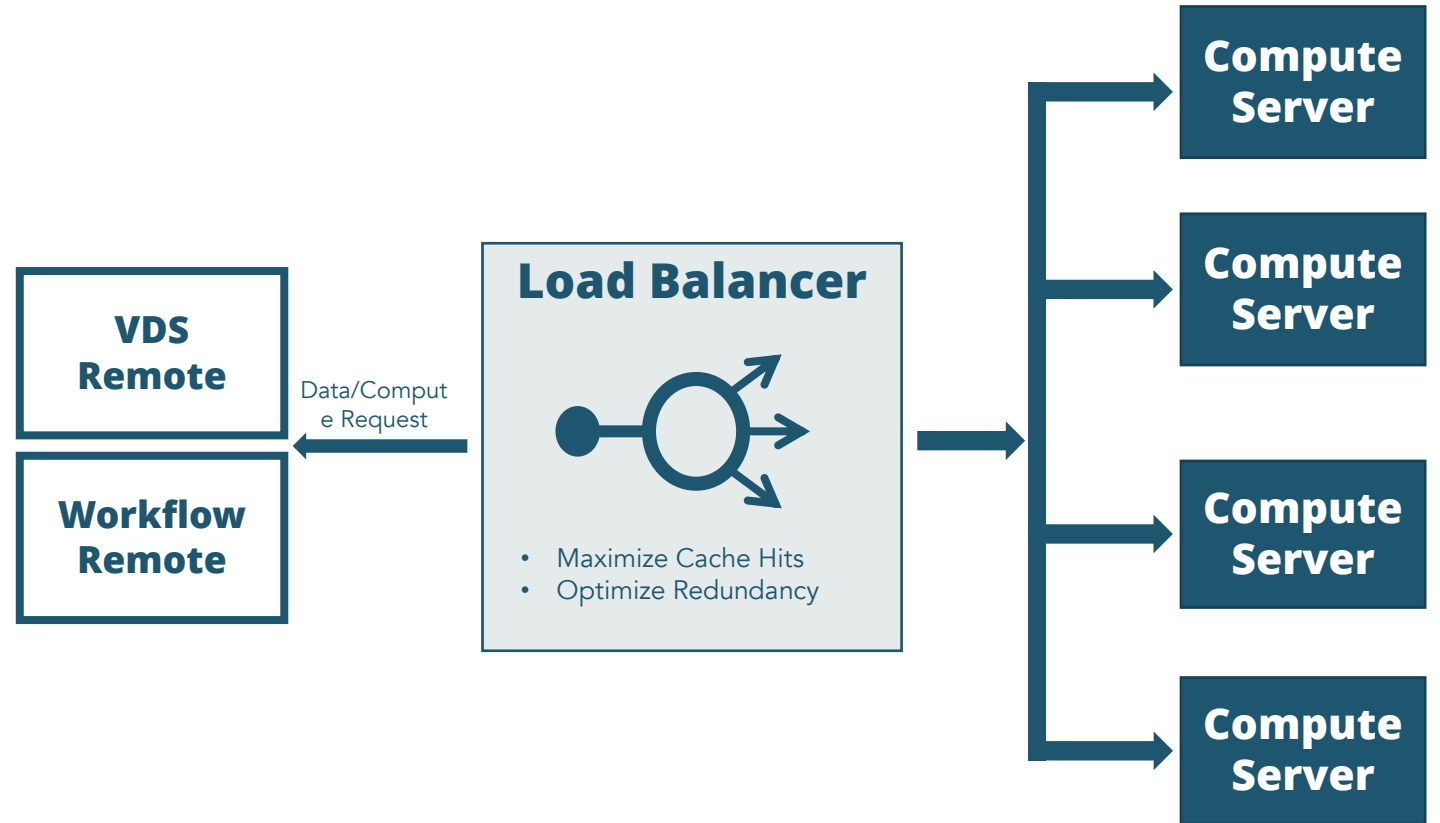
Compute Server

- Each compute node performs workflow compute requests on one or more data blocks.
- Retrieves requested workflow state from workflow state manager and deserializes compute request.
- Stateless other than performing computation for requested data blocks.



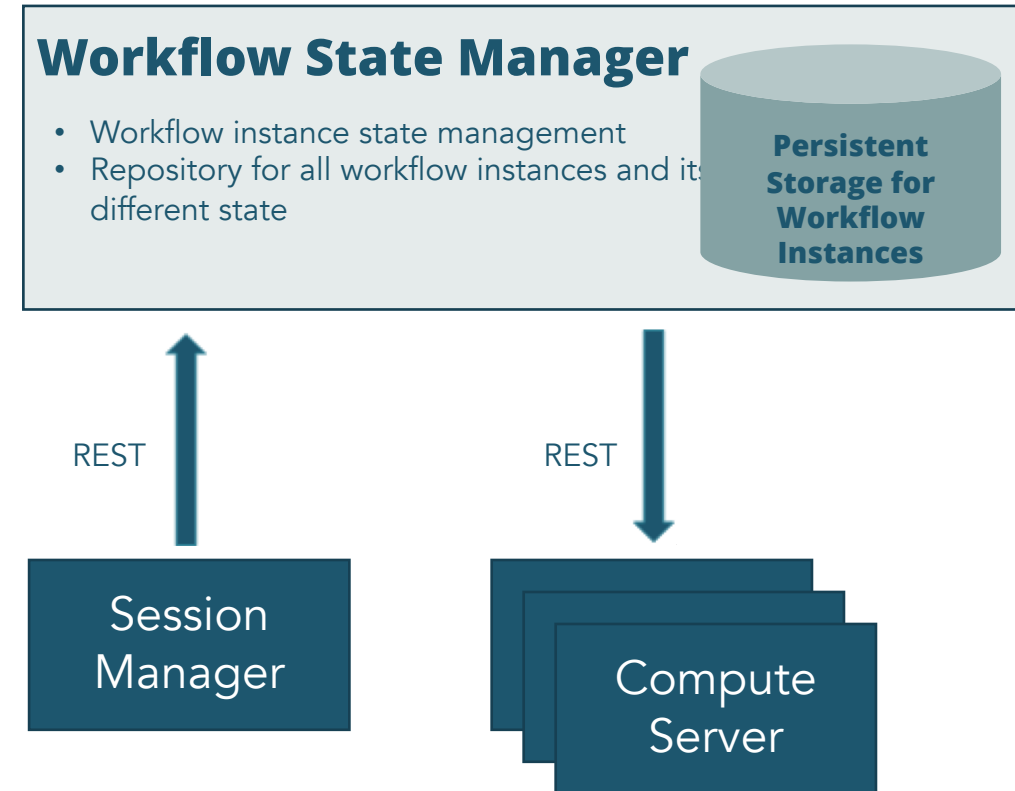
Load Balancer

- Parallelizes compute request to available compute nodes.
- Gathers computed results from different nodes and serves to client.
- Maintains local dictionary of request routing.
- Optimizes performance by routing repeated or geo-spatial dependent compute requests to same compute node.



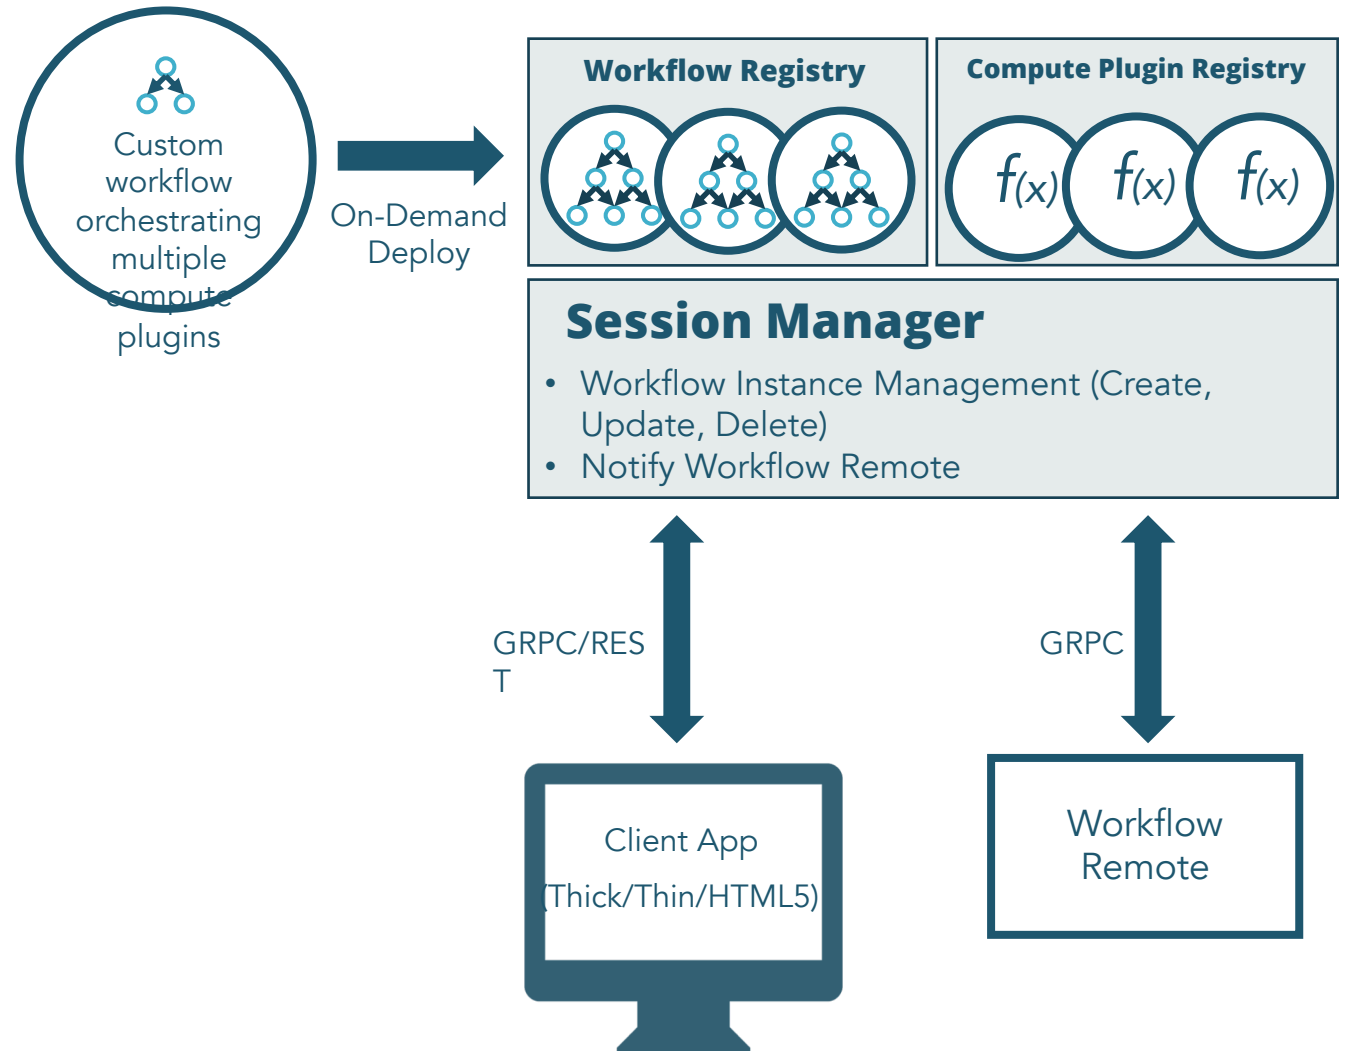
Workflow State Manager

- A single-instance registry of all workflows serialized states.
- All incremental changes to a workflow instance are recorded and persisted.
- Event sourcing from original state of a workflow with intermediate restore points for efficiency.
- History and versioning of computations (traditionally typical of datasets).
- In this ecosystem:
 - *Session Manager* is the *data producer*.
 - *Compute Server* instances are the *data consumers*.
- REST APIs for pulling/pushing records of changes.

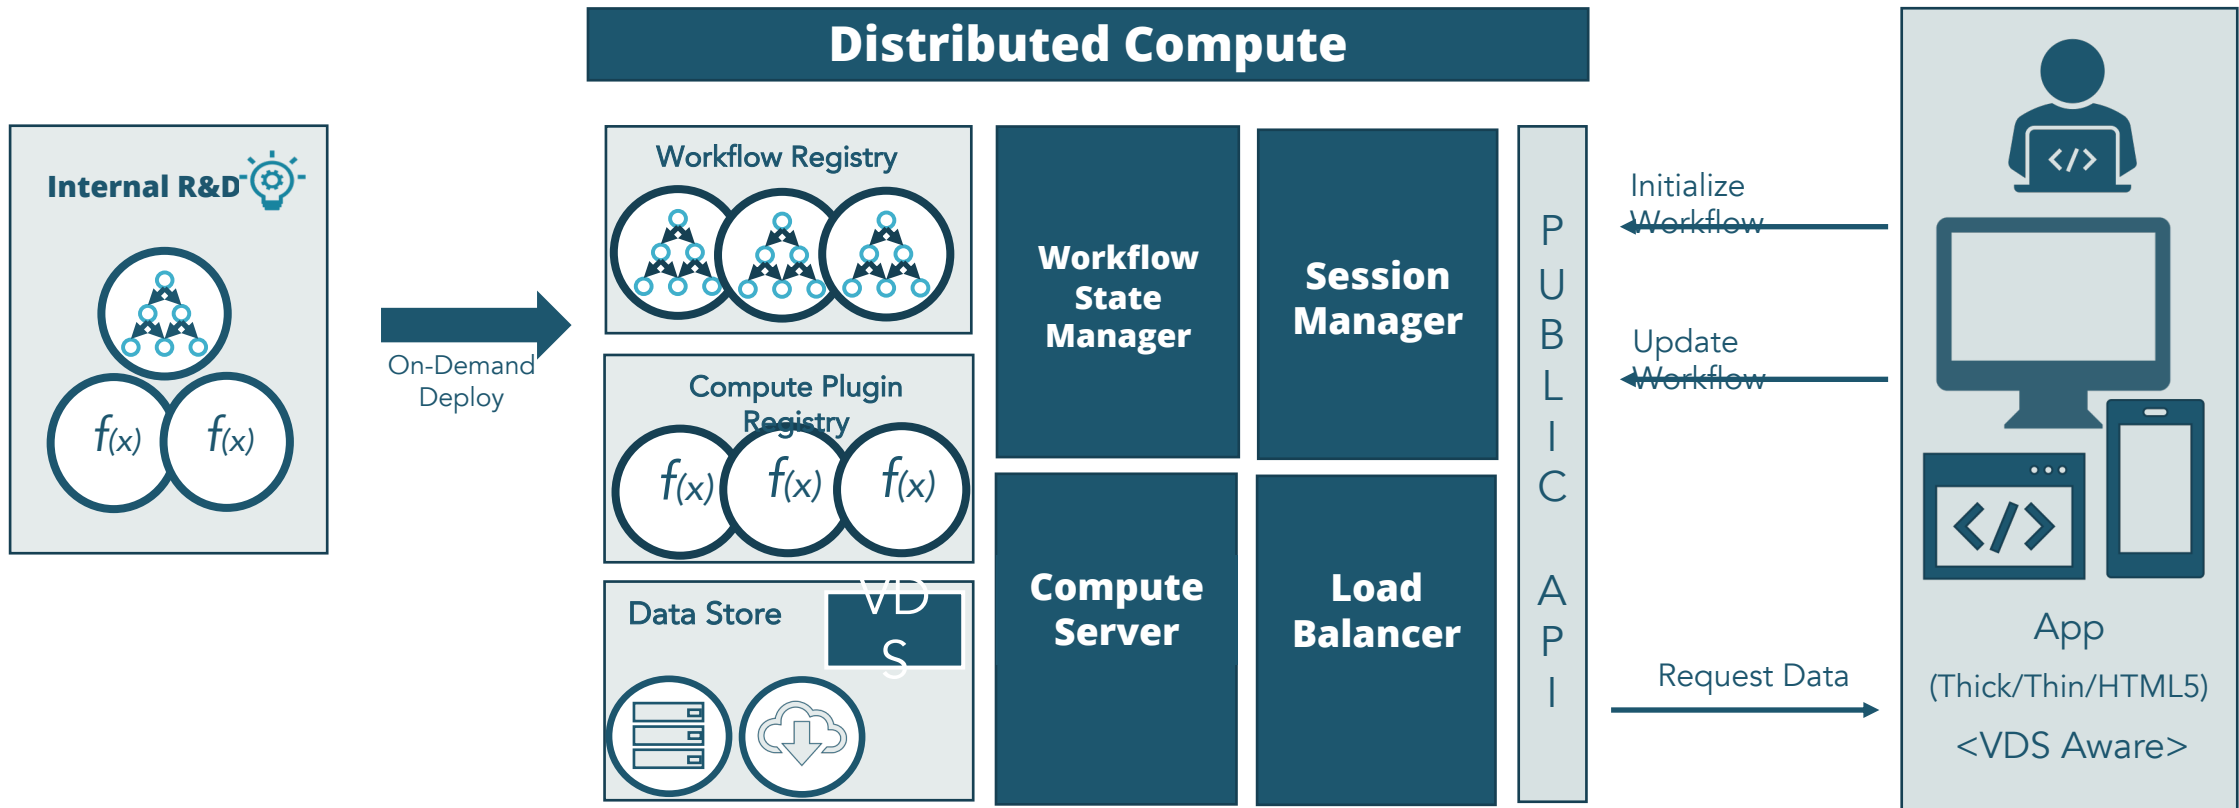


Session Manager

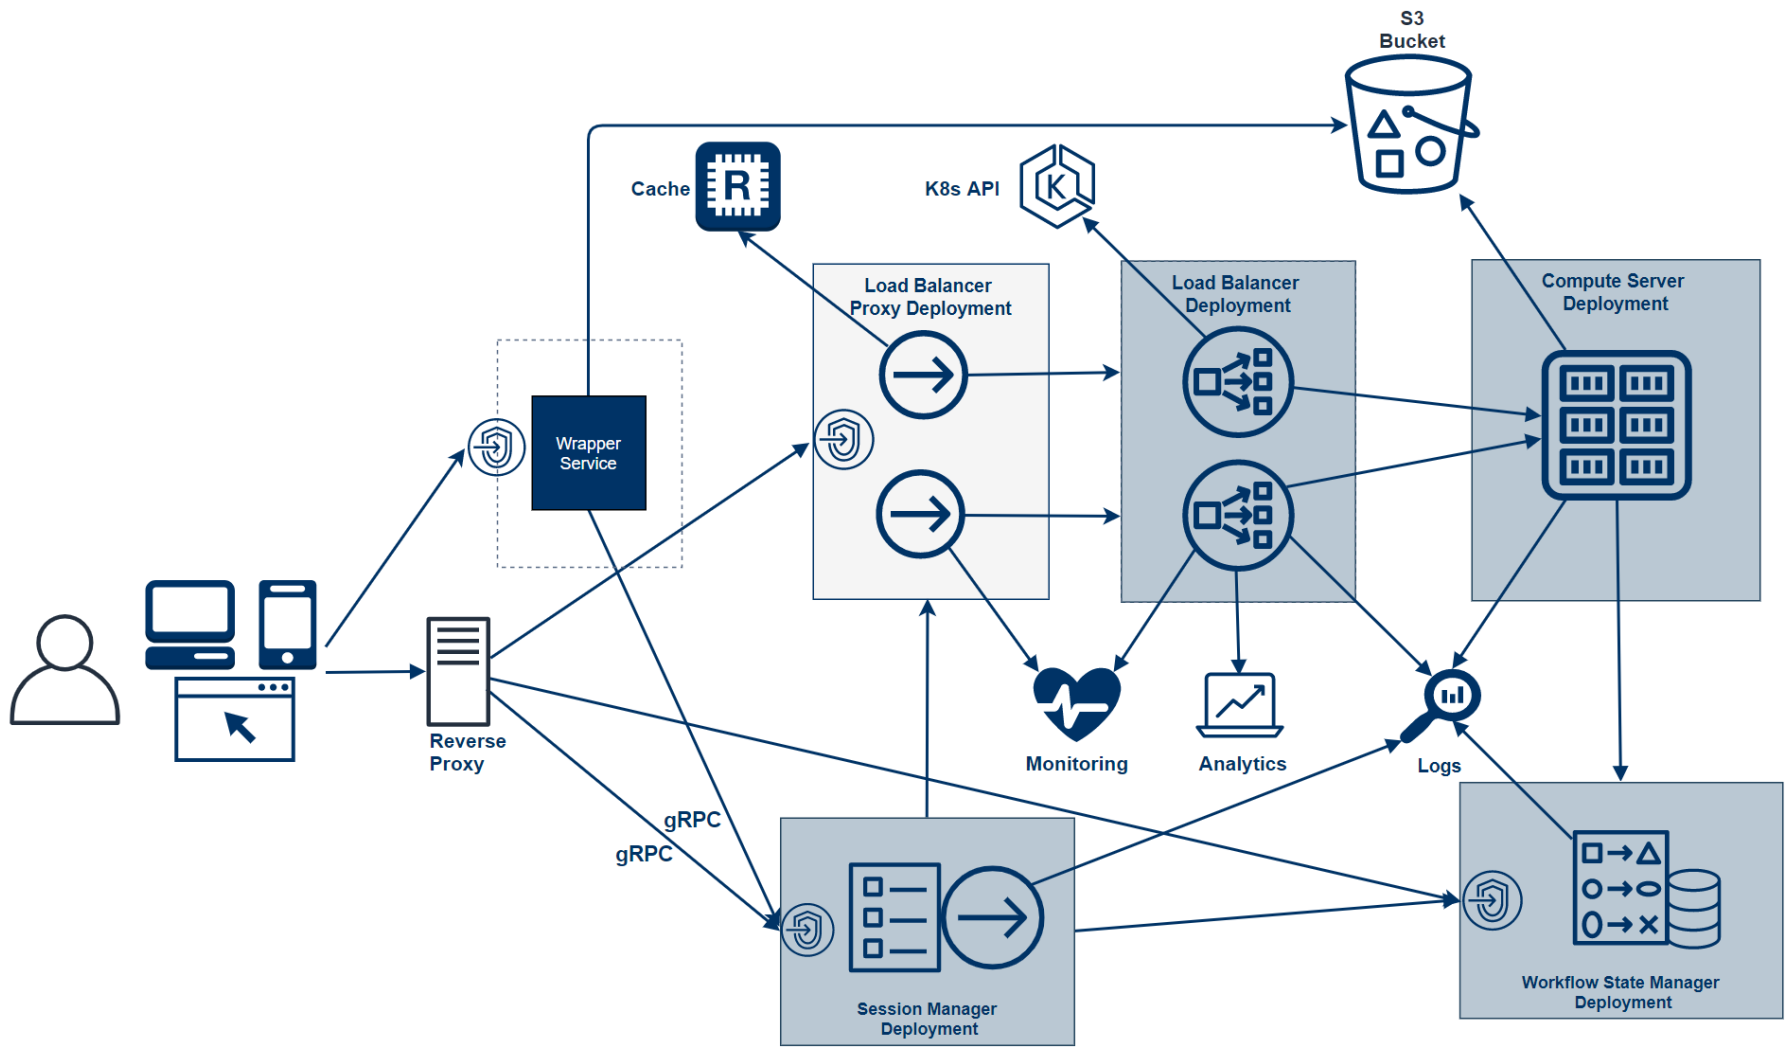
- Discovers new workflows.
- Provides API to search available workflows.
- Provides APIs to create or update workflow instances.
- Coordinates workflow updates with workflow state manager.
- Notifies client application (WorkflowRemote) when corresponding workflow instance is updated.



Distributed Compute – Reference Architecture





Distributed Compute - Cloud Deployment



Scaling and Load Balancing Tests - CPU





Model	c5.4xlarge
CPUs	16
Main Memory	32 GiB
Network Bandwidth	Up to 10 Gbps



Scaling on GPU Nodes



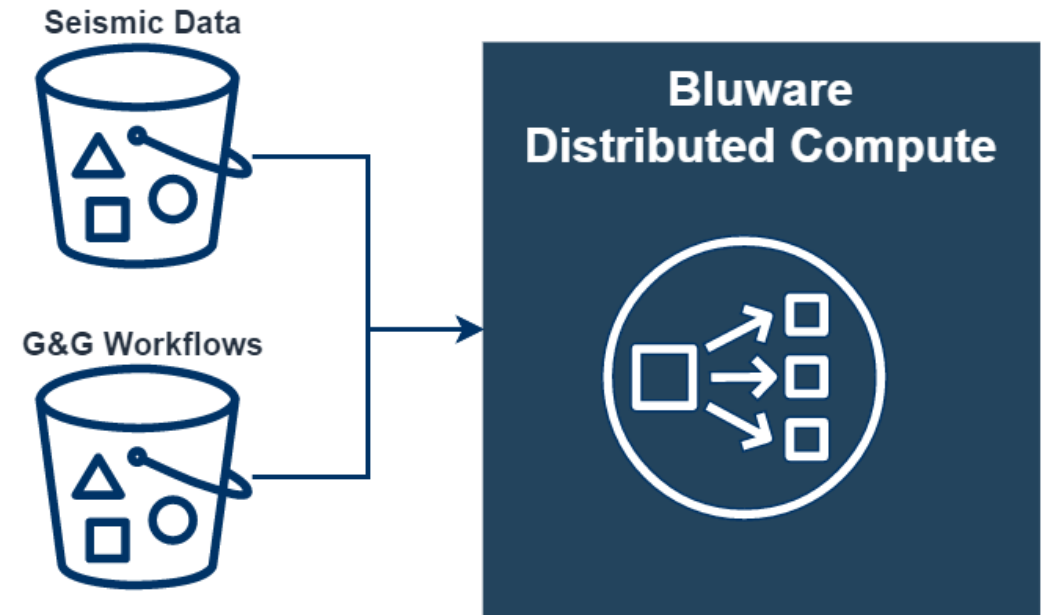
Model	p3.2xlarge
NVIDIA Tesla V100 GPUs	1
GPU Memory	16 GiB
CPUs	8
Main Memory	61 GiB
Network Bandwidth	Up to 10 Gbps



Bluware Distributed Compute



- High Performance Storage
- Optimized and Scalable Compute Engine
- Dynamic Workflows
- Interactive HPC in the Cloud





Tanmoy Palit, *Technical Lead*
Tanmoy.palit@Bluware.com
832.347.9029

Bluware Corp.
16285 Park Ten Place, Suite 300
Houston, TX 77084
713.335.1500

