

Spectrum Scale on IBM Cloud

Spectrum Scale User Group Meeting 2022
SC22 – November 13th, 2022

Piyush Chaudhary (IBM)
HPC & Spectrum Scale Architect – IBM Cloud



IBM Cloud HPC Offerings

- IBM Cloud Catalog BYOL tile offerings with automated provisioning of HPC clusters using
 - **Workload schedulers:** IBM Spectrum Symphony (financial workloads) & IBM Spectrum LSF (non-financial workloads), SLURM (open-source scheduler)
 - **High performance storage:** IBM Spectrum Scale
- Automated deployment, elastic scaling, tools for data management and single vendor tool support
- Hybrid cloud connectivity and secure configurations
- Based on IBM Research benchmarks and current capabilities targeting four workload types

Spectrum Computing Team
(LSF, Symphony)
has moved to IBM Cloud!!

Symphony + Spectrum Scale



Financial Services Sector (FSS)


- Risk Analytics
- Regulatory compliance
- Monte Carlo
- FSS Asset Pricing

LSF + Spectrum Scale




Electronic Design Automation (EDA)

- Optical Proximity Correction (OCP)
- Design Rule Checking (DRC)
- Simulations (like timing analysis)



Life Sciences

- Genomic sequencing
- Drug discovery
- Molecular modeling
- Protein docking

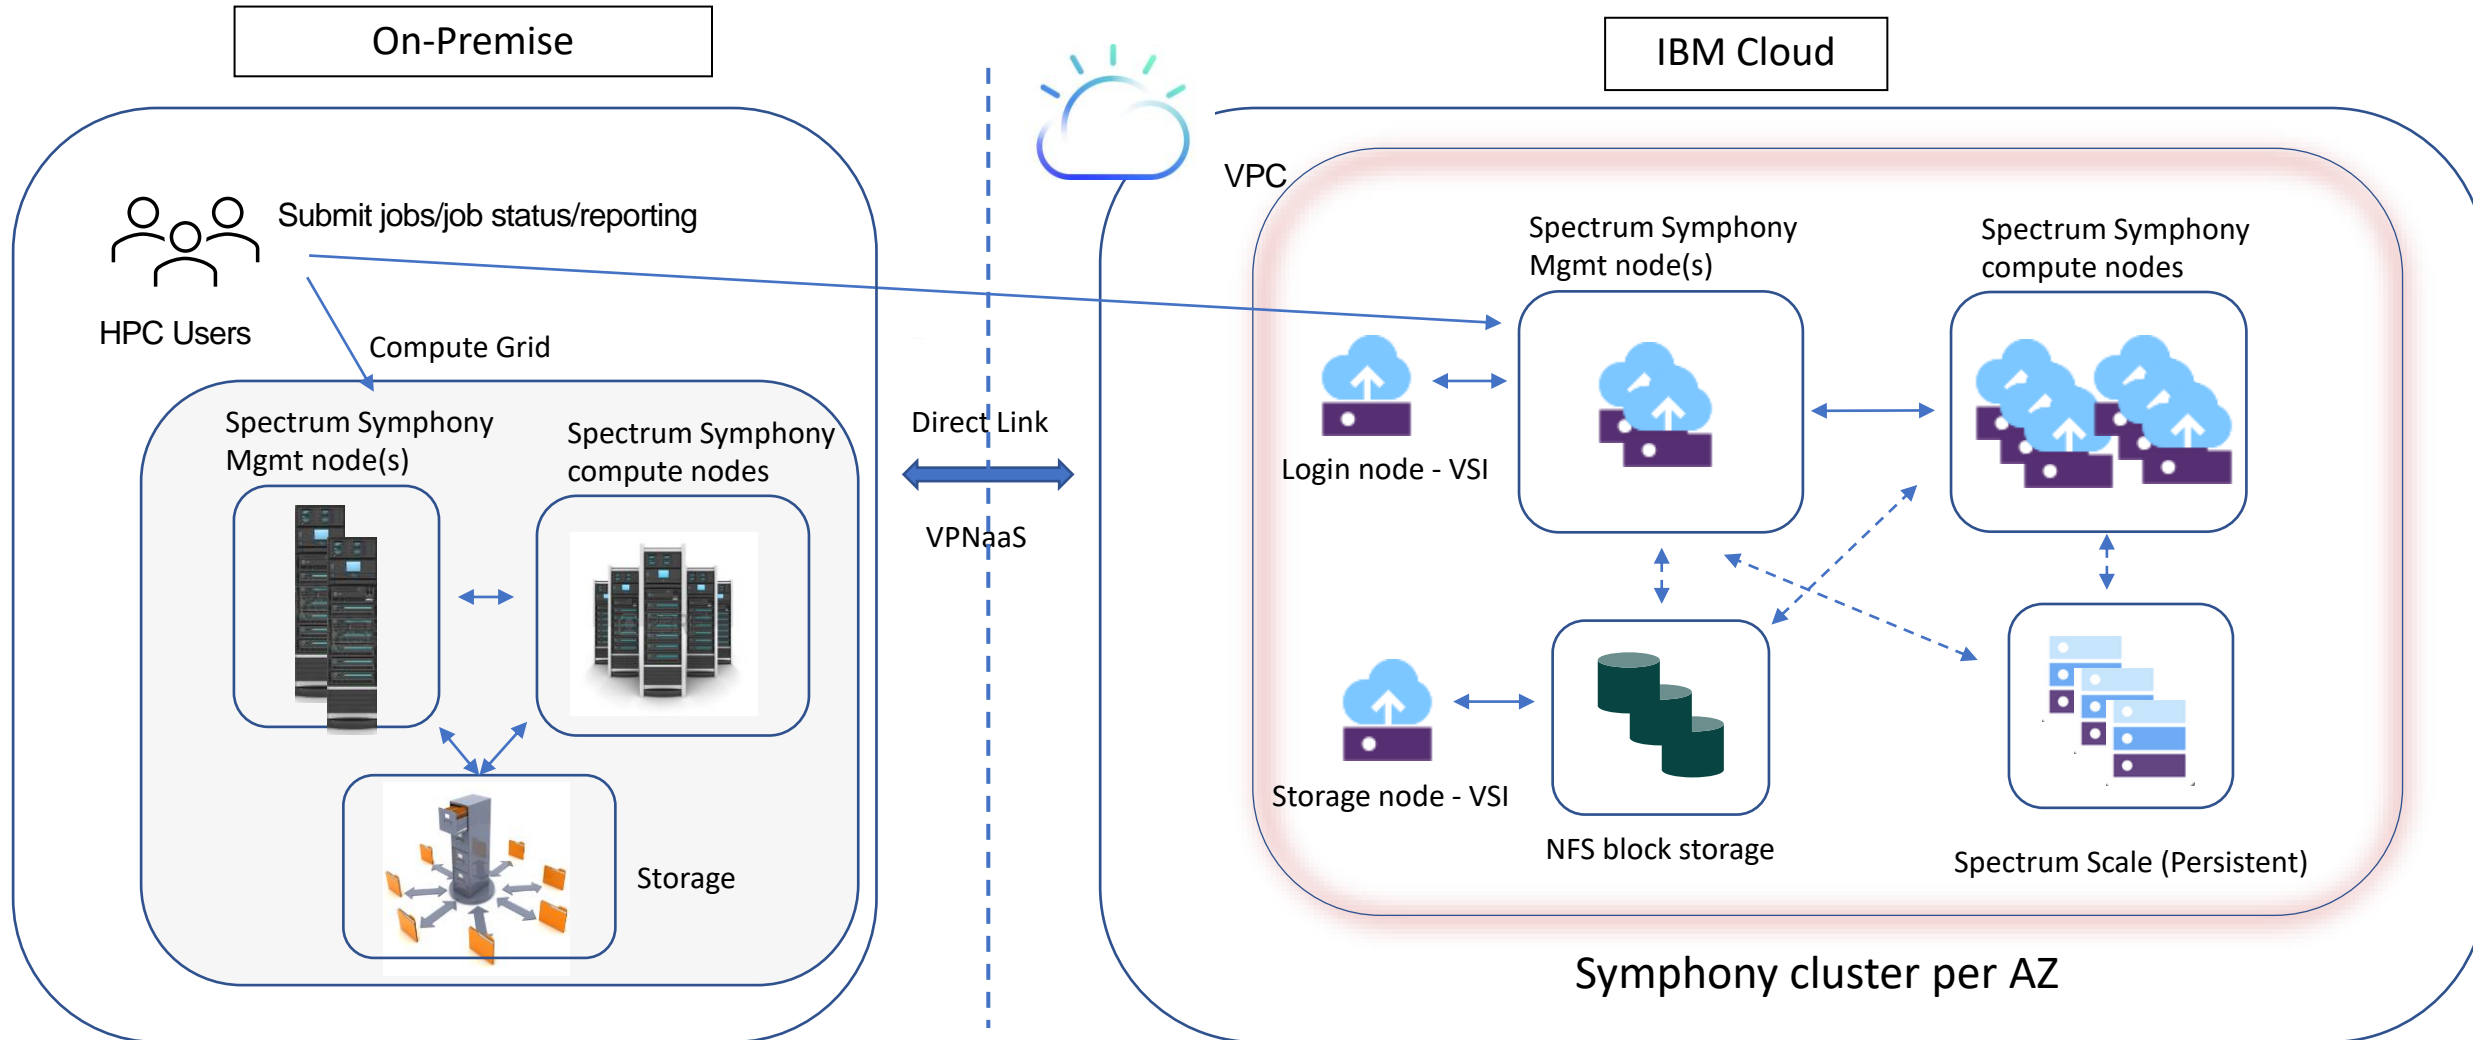


Weather

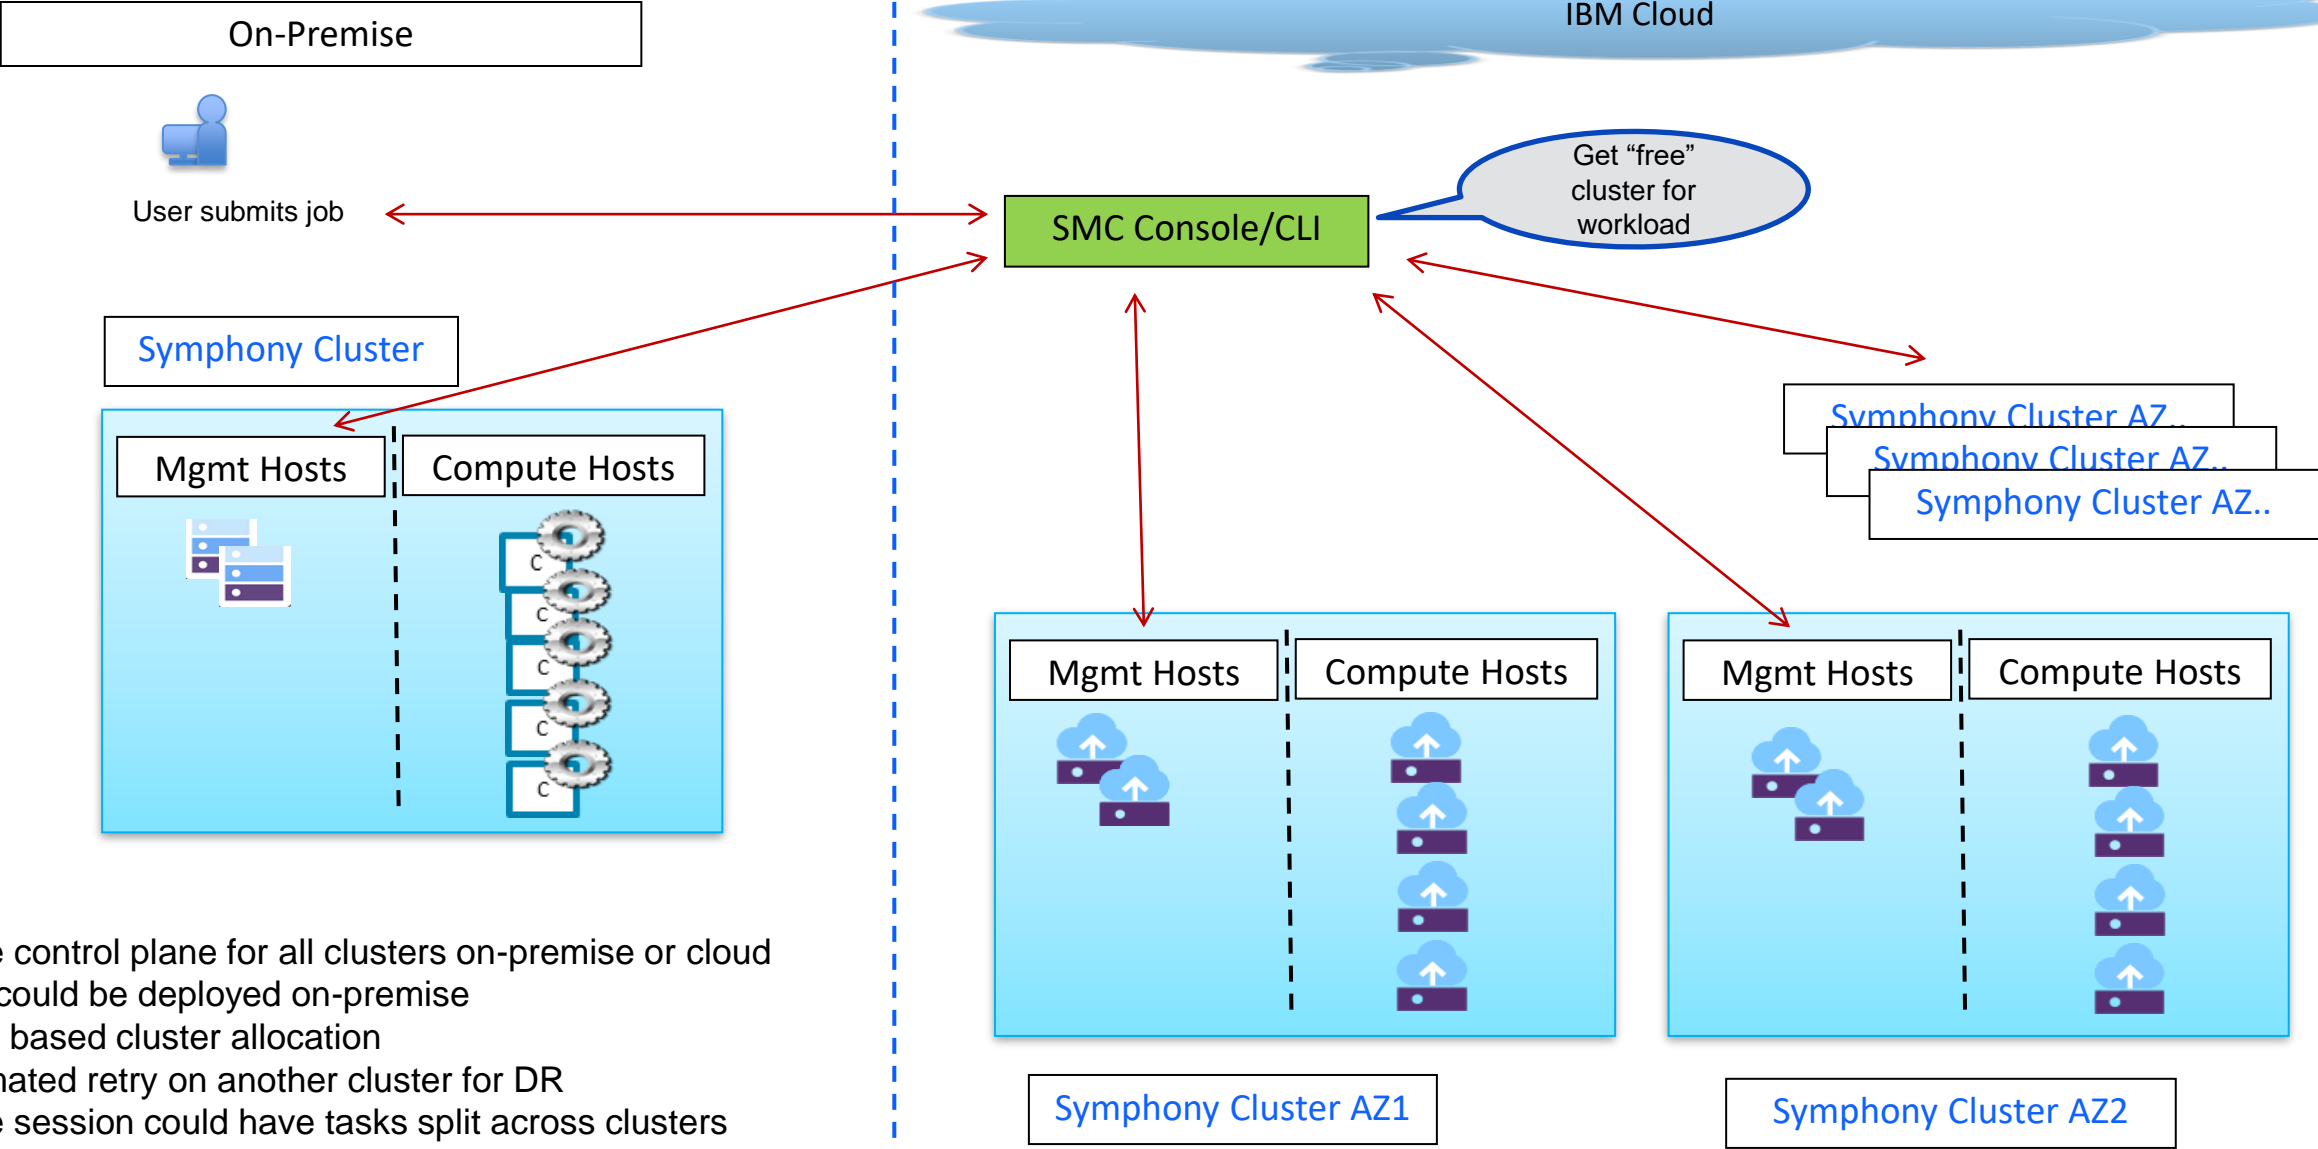
- Climate modeling
- Forecasting and Severe storm prediction
- Ocean Modeling

Reference Architecture

- Hybrid cloud with Direct Link/VPNaaS OR standalone in cloud
- Independent Spectrum Symphony cluster per AZ, leveraging IBM Cloud VPC infrastructure
- Spectrum Scale for data intensive workloads
- BYOL for Spectrum Symphony and Spectrum Scale



IBM Spectrum Symphony Multi-Cluster (SMC)



- Single control plane for all clusters on-premise or cloud
- SMC could be deployed on-premise
- Policy based cluster allocation
- Automated retry on another cluster for DR
- Single session could have tasks split across clusters
- **Easy expansion by adding more compute to a cluster, or a new AZ**

Recent Success Story

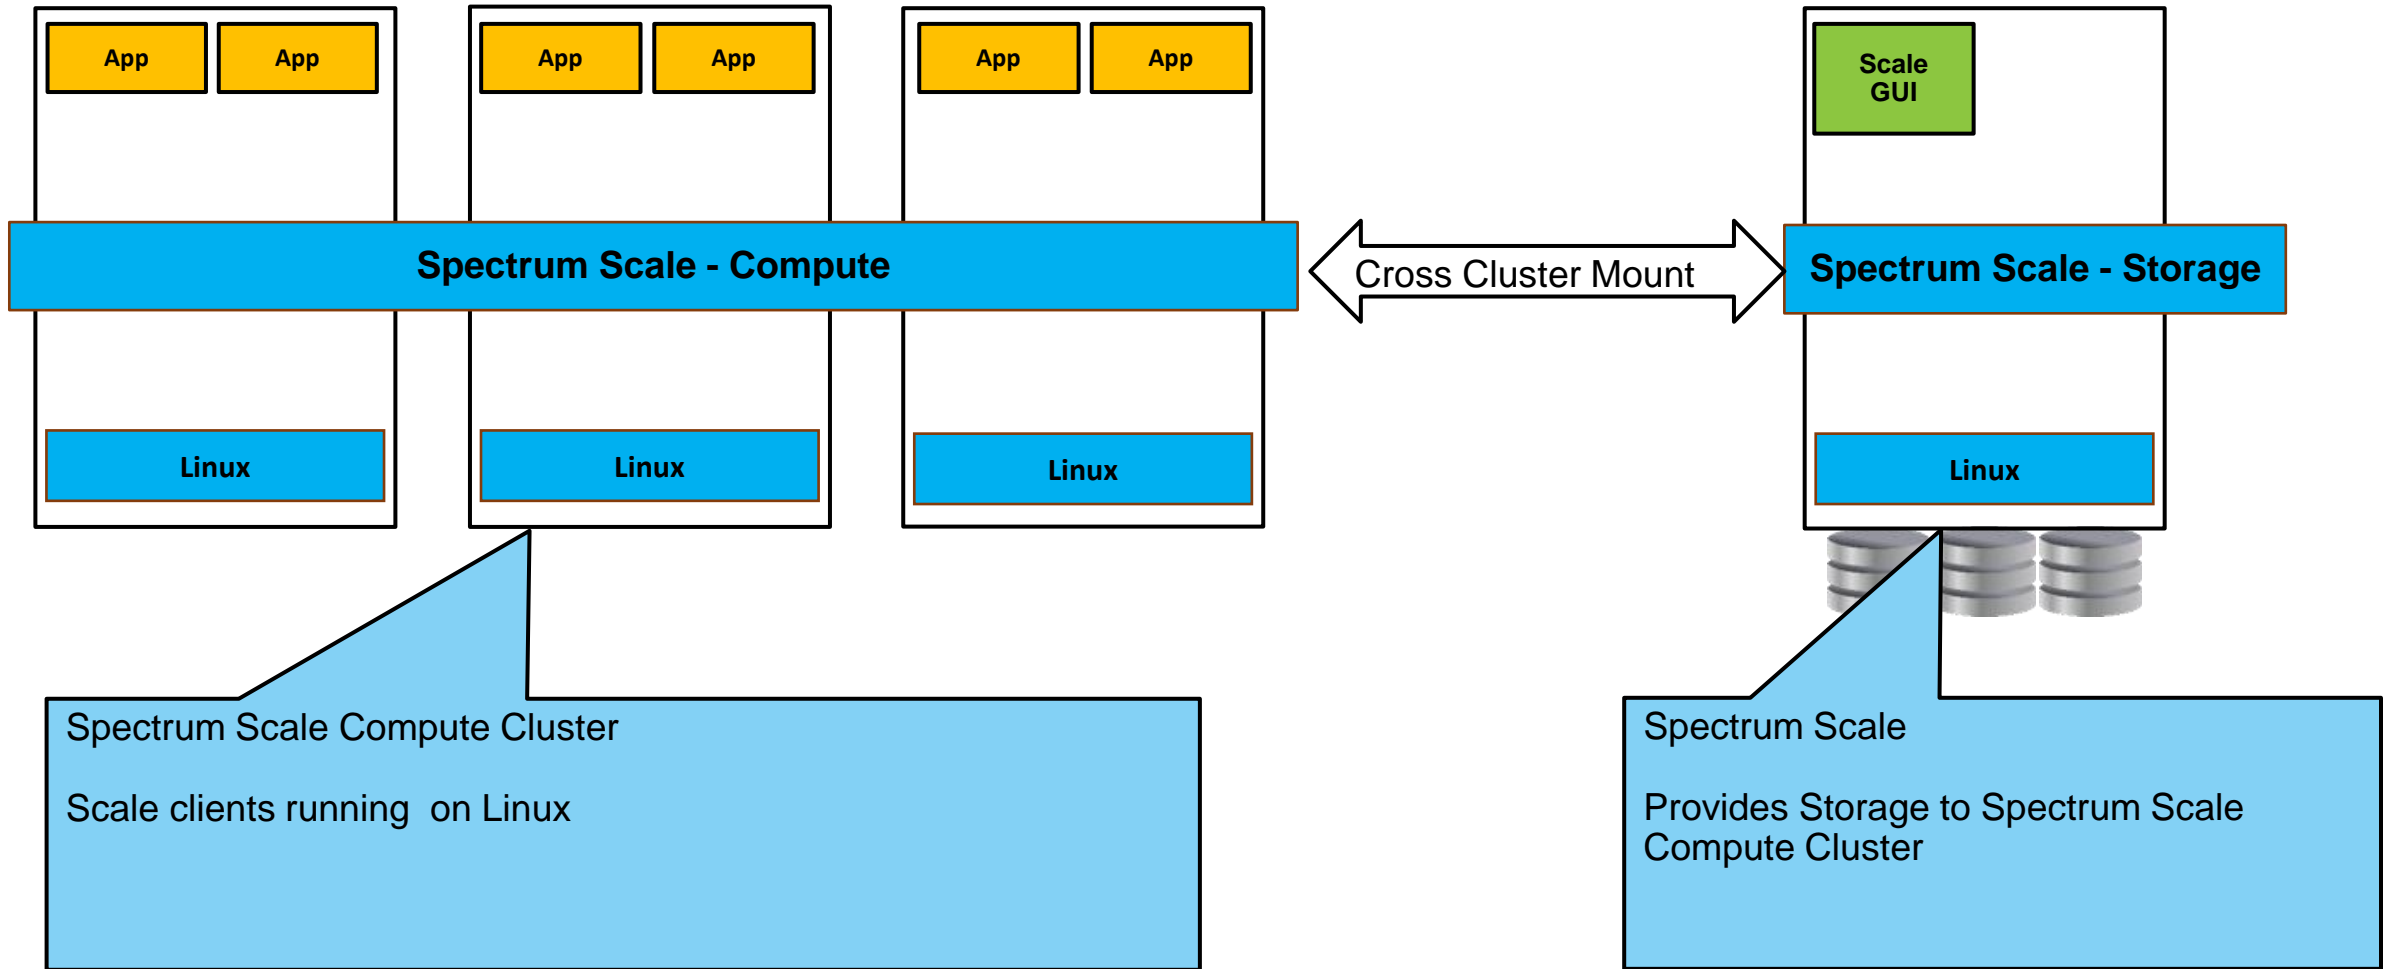


- HPC-as-a-service using managed service wrapper from IBM Consulting, multi year contract
- IBM Spectrum Symphony & IBM Spectrum Scale
- Start with two applications, adding a third & fourth soon
- Starting with 10K vCPU, ramp up to 200K vCPU
- Distributed across 3 Regions in North America, each with 2 AZ's

Key Factors for Win

- Deep HPC Expertise, only IBM knows Symphony & Spectrum Scale solutions
- ONE IBM value for customer, deep collaboration across business units, solutioned all issues effectively
- Customer can focus on critical business application development
- PoC success with first application, lead to second, third and now fourth application wins!

Spectrum Scale in IBM Cloud - Architecture



Spectrum Scale Tile in IBM Catalog

Catalog /

IBM Spectrum Scale

IBM Spectrum Scale is a high-performance, highly available, clustered file system and associated management

storage_type

Type
Terraform

Provider
IBM

Last updated
11/12/2022

Category
Storage

Related links
[Docs](#)
[Get help](#)
[Readme](#)

Content source URL ⓘ

<https://cm.global...>

Select version

2.1.1

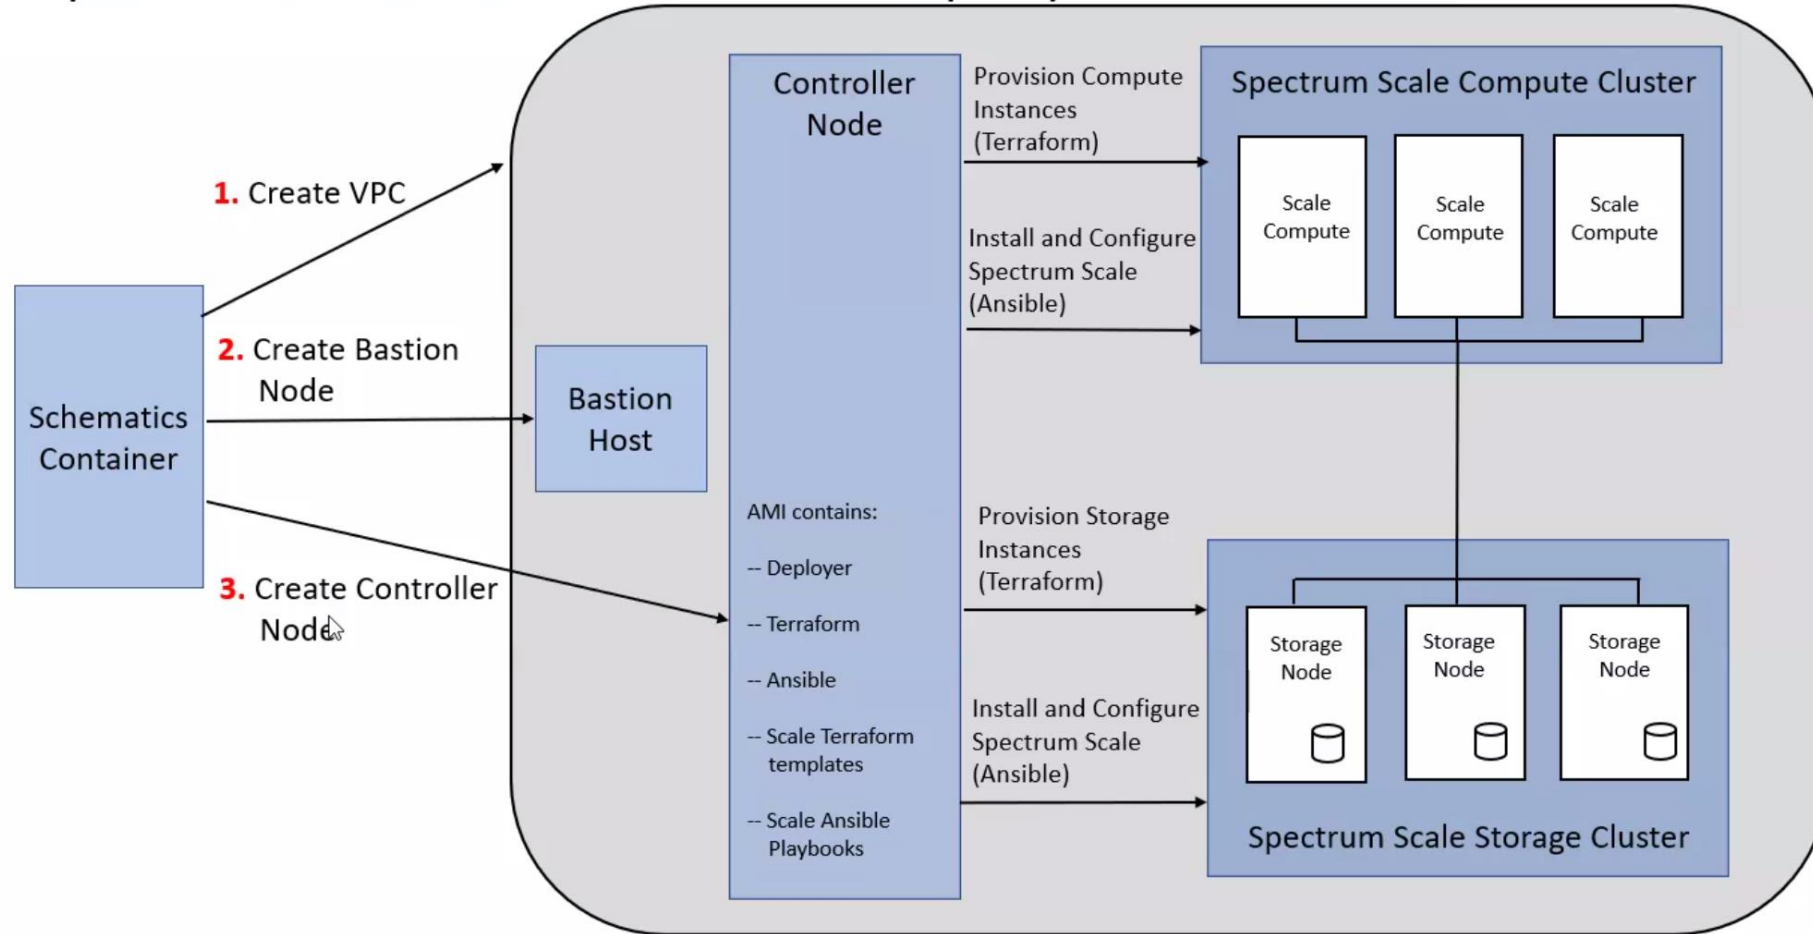
Select the Spectrum Scale file system deployment method. Note: The Spectrum Scale scratch and **evaluation type** deploys the Spectrum Scale file system on virtual server instances, and the persistent type deploys the Spectrum Scale file system on bare metal servers. The persistent Spectrum Scale storage feature is a beta feature that is available for prototyping and testing purposes. There are no warranties, SLAs, or support provided for persistent storage and it is not intended for production use.

Interact to collapse Tile




Spectrum Scale on IBM Cloud - Deployment

Spectrum Scale Controller Deployment Model




Spectrum Scale – Storage Cluster



Scratch
VSIs with Instance
Storage

+
Evaluation



Persistent
Bare Metal Nodes with
NVMe Drives




Additional Capabilities (Future)


- Erasure Code Edition (ECE)
- Cross AZ replication
- Async replication to S3
- Protocol integration
- AFM connection to on-prem
- Backup/Restore
- Archive



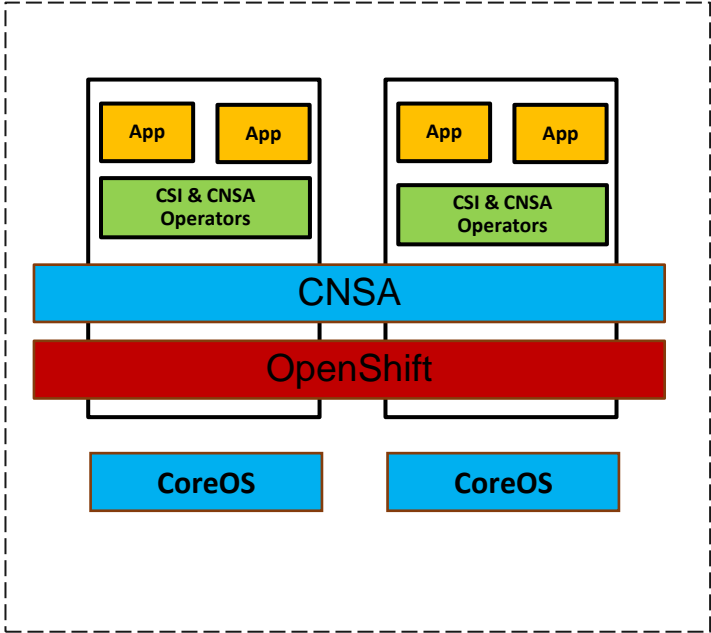
Spectrum Scale – Compute Cluster



Native Spectrum Scale Access



High Performance NFS Access



Container Native Access

Performance Preview

Cluster Runtime Evaluation

- 10 storage nodes
 - 80 NVMe drives
 - 256 TB
 - 100 Gbps network / storage node
- 64 compute nodes
 - cx2-16x32 instance profile
 - 512 physical cores
 - 24 Gbps network
- Spectrum Scale 5.1.4

Performance Preview

Cluster Deployment Settings

- Storage nodes
 - GPFS replication factor: 2
 - MTU 9000
 - Ethernet device combined queues: 16
- Compute nodes
 - MTU 9000
- Performance testing conducted with out-of-the-box configuration

Performance Preview

File Size & Total Files	Operation	Performance	Procs per Compute Node	Threads per Process	Request size
64 files 1 TiB	Write sequential	37 GiB/sec	1	8	4MiB
	Read sequential	123 GiB/sec	1	12	4MiB
	Write random	876,002 IOPS	1	136	4KiB
	Read random	5,601,396 IOPS	1	80	4KiB
128 Files 512 GiB	Write sequential	38 GiB/sec	2	4	4MiB
	Read sequential	125 GiB/sec	2	4	4MiB
	Write random	873,046 IOPS	2	68	4KiB
	Read random	5,605,109 IOPS	2	40	4KiB



Spectrum Scale – Roadmap

(Subject to Change)

Feature Enhancements

- Erasure Code Edition (ECE)
- Cross AZ replication
- Async replication to S3
- Protocol integration
- AFM connection to on-prem
- Backup/Restore
- Archive

Use Cases

- FSS
- Life Sciences / Healthcare
- Electronic Design Automation (EDA)
- Weather / Climate
- Automotive
- Analytics & AI
- Industrial Manufacturing

Thank you for using
IBM Spectrum Scale on IBM Cloud!