


IBM at SC17 Spectrum Scale User's Group Meeting

Transparent Cloud Tiering



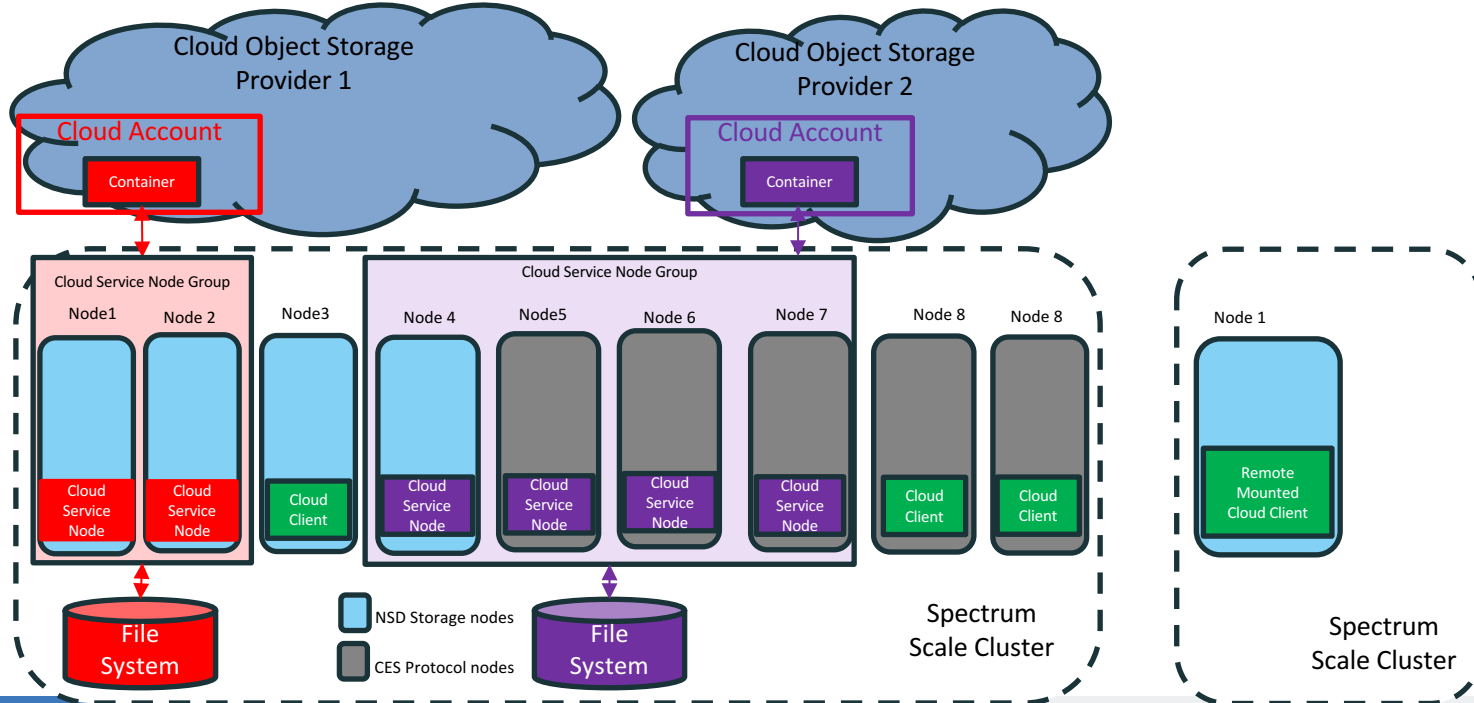
Spectrum Scale 4.2.3 Transparent Cloud Tiering scaling

- **Node Groups** – Up to 4 node groups of 4 nodes each per cluster
- **Multi-cluster support:** No support for multiple remote cluster access from a given cluster
- **File system support** – 1 file system per Cloud Service Node group
- **Cloud account support** – Single cloud storage tier: 1 Cloud Account per filesystem with one cloud storage access point
- **Filesystem scaling:** One Container per filesystem



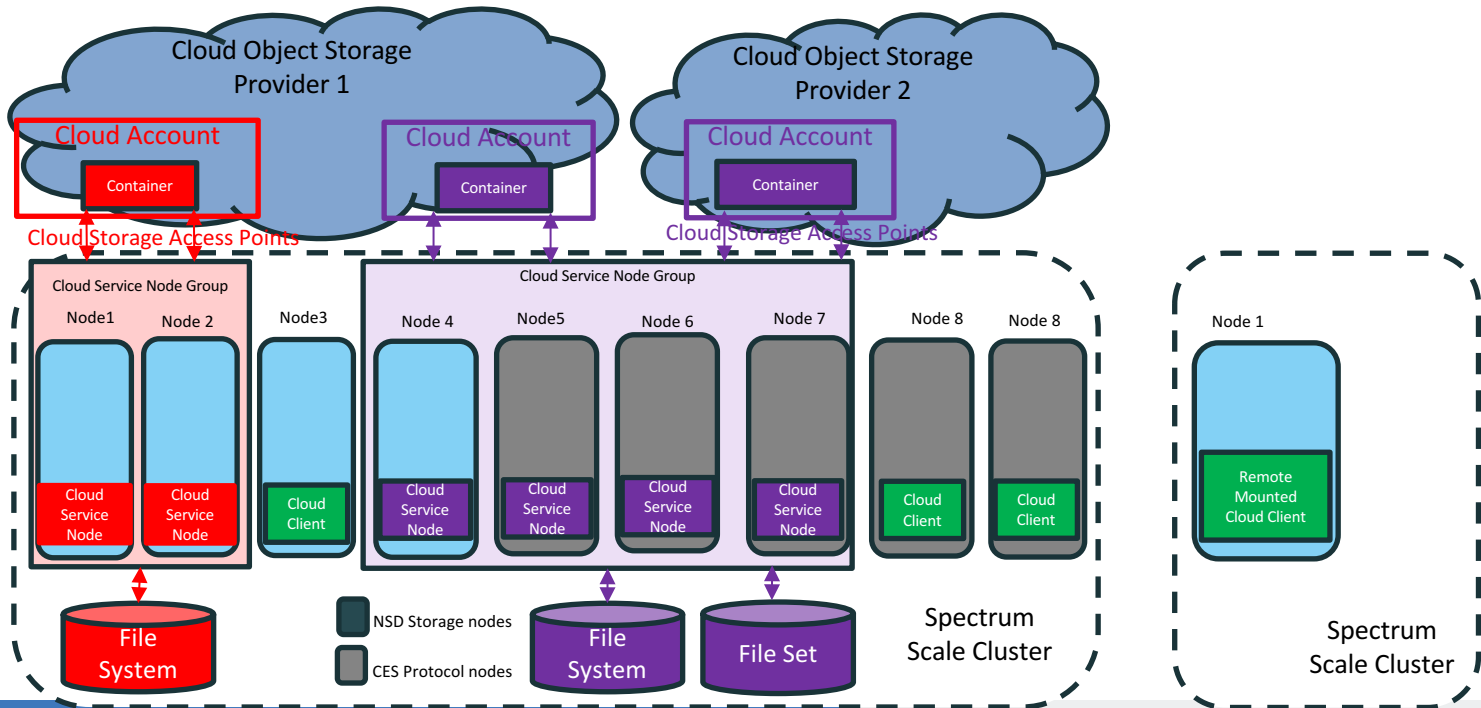
Spectrum Scale 5.0 Multi-cluster scaling

- **Node Groups** – Up to 4 node groups of 4 nodes each per cluster
- **Multi-cluster support:** - Multiple remote cluster access from a given cluster




Spectrum Scale 5.0 Multiple Cloud Account Support

- **Node Groups** – Up to 4 node groups of 4 nodes each per cluster
- **Multi-cluster support:** Multiple remote cluster access from a given cluster
- **File system support** – Multiple filesystems or file sets per Cloud Service Node group (on the order of dozens)
- **Cloud accounts support** – Two Cloud Targets or Tiers: Up to 2 Cloud Accounts per filesystem, each with multiple cloud storage access points ([Regions](#), [Accessor URLs](#))



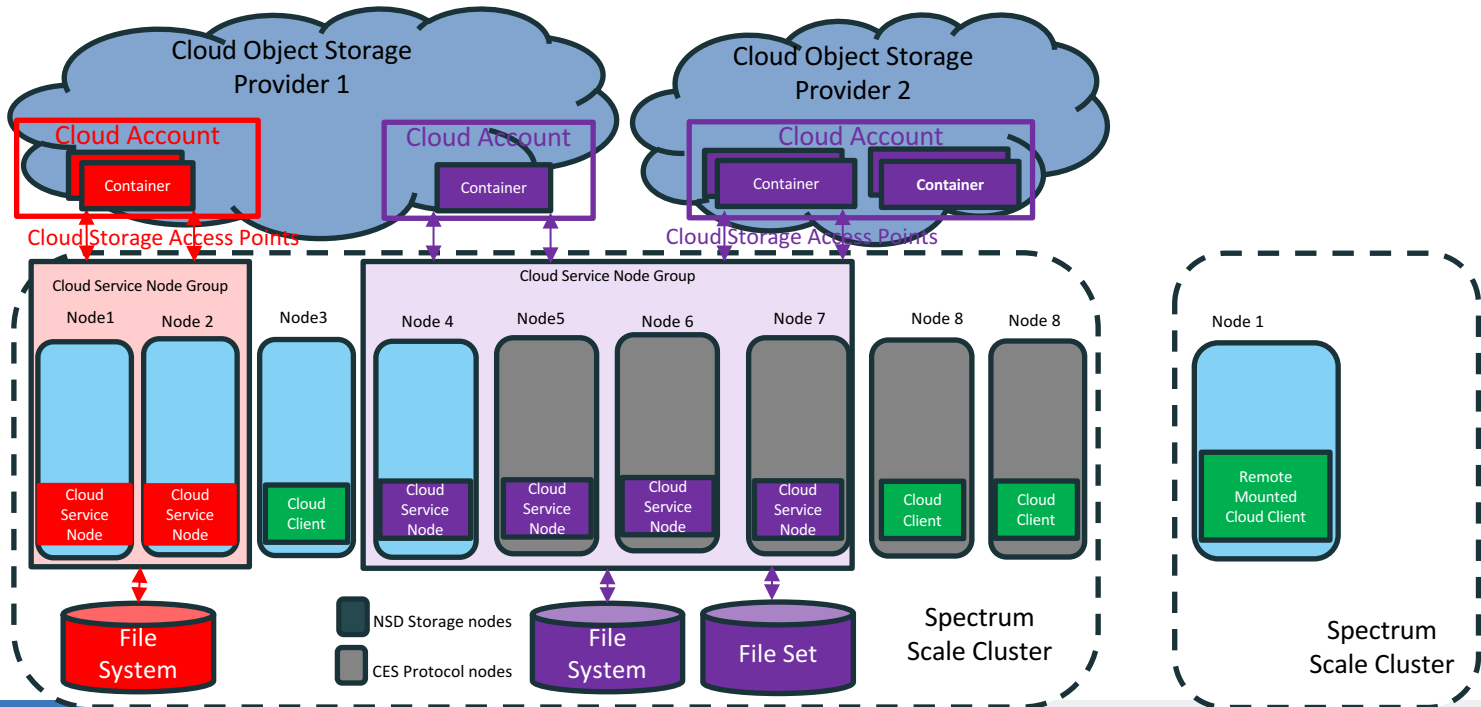
Spectrum Scale 5.0 Multiple Cloud Account Support

- **Node Groups** – Up to 4 node groups of 4 nodes each per cluster
- **Multi-cluster support:** Multiple remote cluster access from a given cluster
- **File system support** – Multiple filesystems or file sets per Cloud Service Node group (on the order of dozens)
- **Cloud accounts support** – Two Cloud Targets or Tiers: Up to 2 Cloud Accounts per filesystem, each with multiple cloud storage access points ([Regions](#), [Accessor URLs](#))



Spectrum Scale 5.0 Container Spill-over

- **Node Groups** – Up to 4 node groups of 4 nodes each per cluster
- **Multi-cluster support:** Multiple remote cluster access from a given cluster
- **File system support** – Multiple filesystems or file sets per Cloud Service Node group (on the order of dozens)
- **Cloud accounts support** – Two Cloud Targets or Tiers: Up to 2 Cloud Accounts per filesystem, each with multiple access points (Regions, Accessor URLs)
- **Filesystem scaling:** Multiple container support using container spill-over allows for supporting a large number of files per filesystem or file set
 - Recommended: scaling to 100 Million files per each container, with ability to spill-over to additional containers as needed



Spectrum Scale 5.0 Key Scaling Points

- Insure hardware will meet scaling requirements
 - Multiple nodes – the more cores per node, the better
 - One or more dedicated 10 Gig connections between each node and cloud storage
 - Scale metadata and TCT Cloud DB should in flash
- Follow these configuration best practices
 - Container spill-over at a reasonable number of files
 - Keep up on maintenance activities (deletes, reconciles, cloud DB backup)
 - Tune Spectrum Scale parameters as outlined TCT pubs
- Optimize for what object storage does well
 - Scale out – use lots of threads
 - Use bigger size files where possible
 - Leverage vaults / containers to help with quota support
- File scalability usually depends on one of these two key factors:
 - Peak and steady state bandwidth for node group – Typical numbers for decent hardware might 1-2 Gigabyte /second
 - inode Service Restore time – Typical numbers might be 3 weeks / billion files

Backup