ORACLE

Fastest File Server in Public Cloud using IBM Spectrum Scale on Oracle Cloud Infrastructure

- Pinkesh Valdria

  Principal Solutions Architect - Oracle Cloud

  + in collaboration with Re-Store LLC (Michael Sedlmayer + Hideaki Kikuchi)

Oracle Cloud Infrastructure

Oracle Cloud Infrastructure is a set of complementary public cloud services that enable you to build and run a wide range of applications and services in a highly available hosted environment.

Key Resources	Performance	Price	Reliability
Compute	Most powerful instances available	Up to 52% less than the competition	The only Compute Manageability SLA
Networking	Up to 2x25 Gbps from the instance No oversubscription	Lowest price for getting or serving your data	The only Networking Performance SLA
Block Storage	Superior IOPS/GB Superior IOPS/instance	Lowest price for guaranteed performance (up to 7,900% less)	The only Block Storage Performance SLA

Oracle Cloud Infrastructure Global Footprint

End of CY2020: 36 Oracle Regions

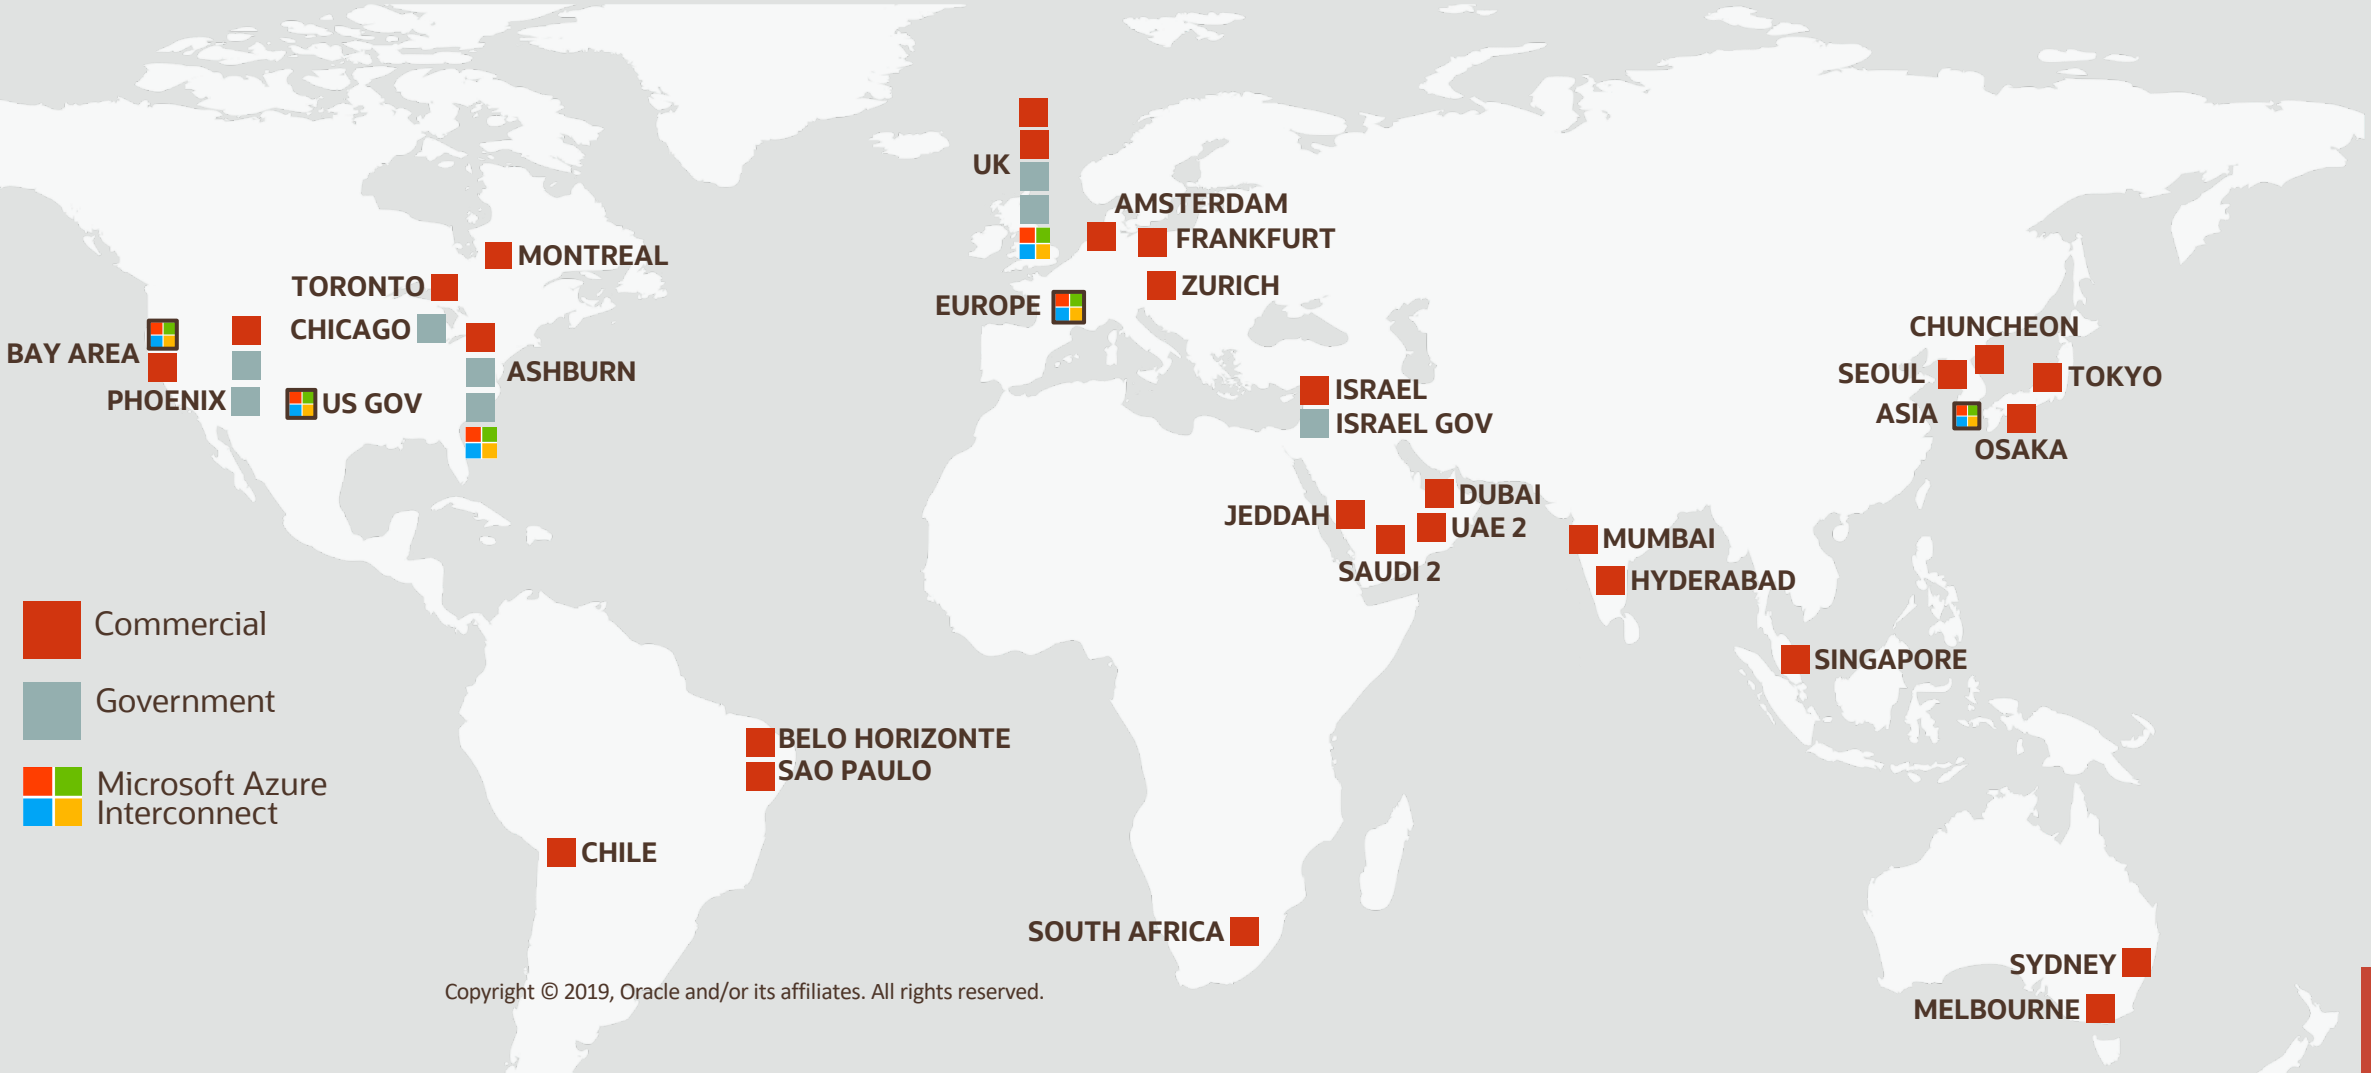

Background

IBM Spectrum Scale (GPFS)

IBM Spectrum Scale™ is a high-performance, highly available, clustered file system and associated management software. IBM Spectrum Scale can scale in several dimensions, including performance, capacity, and number of nodes or instances that can mount the file system.

Use cases

- High performance computing (HPC) workloads
- Analytics/Modeling/BI
- Data intensive applications & workflows – like Audio/Video storage & editing
- SAS Grid Computing

Why Spectrum Scale?

Customer wants

- SAS Grid applications require throughput of 100MB/s/physical core
- NFS challenges
  - Need higher throughput
  - metadata lookup takes too long when millions of files are stored in same directory

Spectrum Scale Architectures on Oracle Cloud

Architectures supported on Oracle Cloud

- Network Shared Disk (NSD) architecture
- Direct Attached Disk architecture (SAN type)
- Erasure Code Edition
- Shared Nothing architecture

Architecture

- Build using Bare Metal Standard Compute shapes and OCI Block Volumes
- **Multi-attach Block Volumes** to multiple compute nodes (Only supported on OCI)
- 25 Gbps network interface to Block storage
- Uses “building block” approach – easy to scale up and manage

Network Shared Disk (NSD) Server Model – Single AD

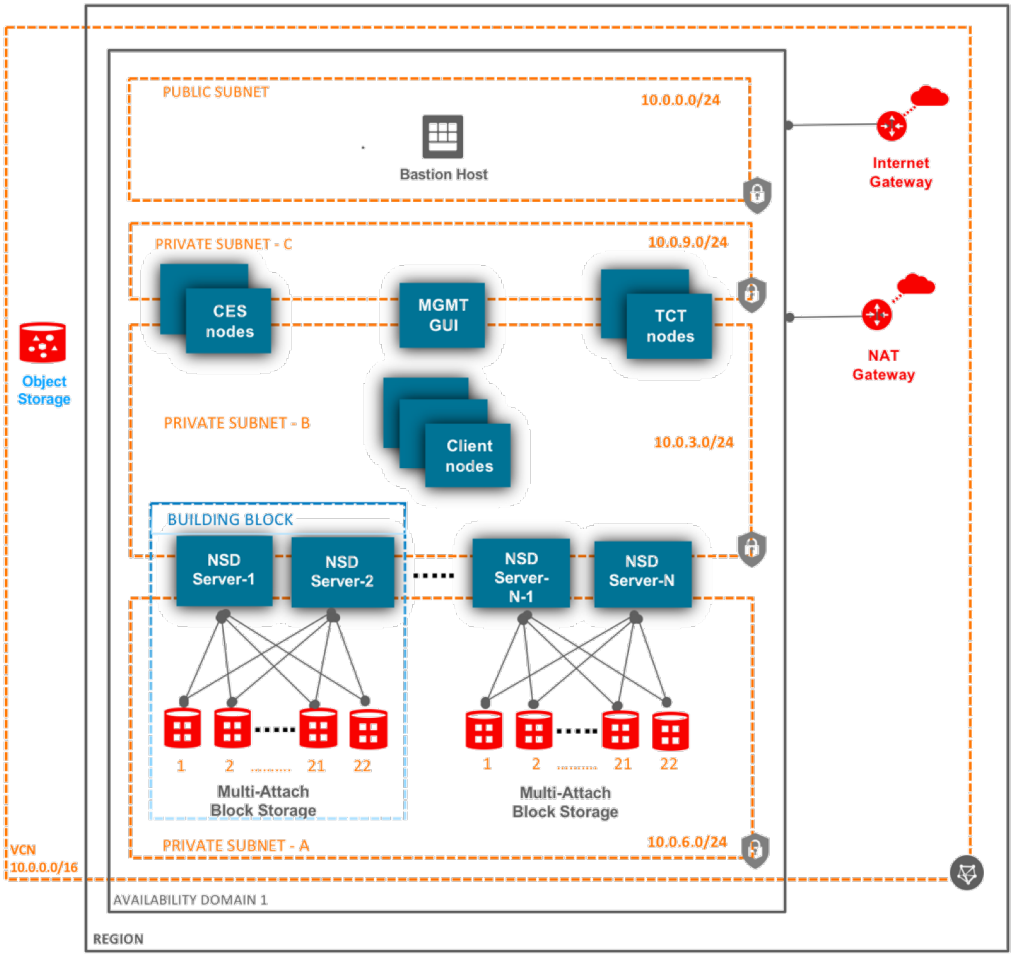

Architecture

- Build using Bare Metal Standard Compute shapes and OCI Block Volumes
- **Multi-attach Block Volumes** to multiple compute nodes (Only supported on OCI)
- 25 Gbps network interface to Block storage
- Uses “building block” approach – easy to scale up and manage

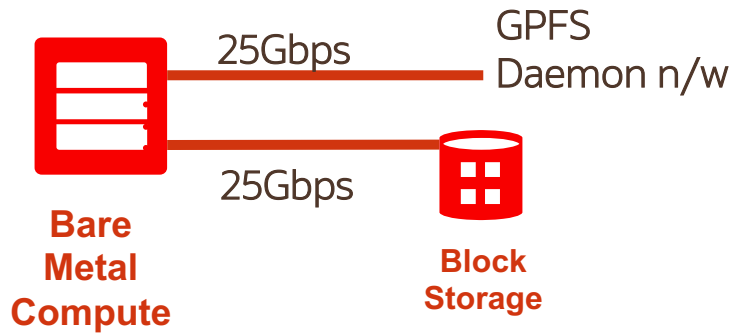

Performance – IOR Benchmark

- On as **small as two nodes**, IBM Spectrum Scale on Oracle Cloud Infrastructure provides over **5 GB/s throughput**
- By adding more building block, the throughput **scales almost linearly**
- Achieve **58 GB/s for read** using just 11 building blocks (22 servers)

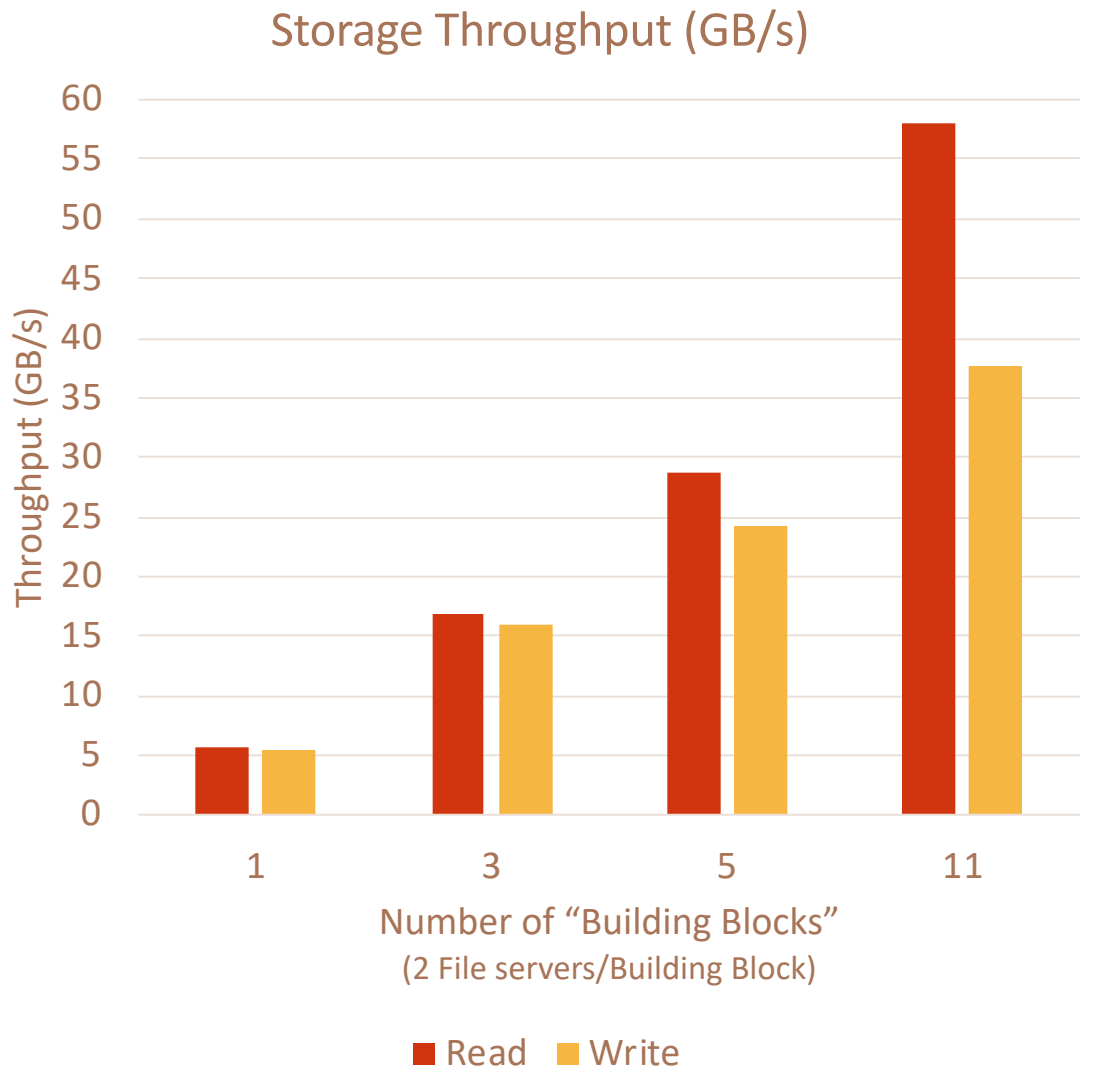

Cost

As you increase the amount of storage, the proportional compute costs decrease.

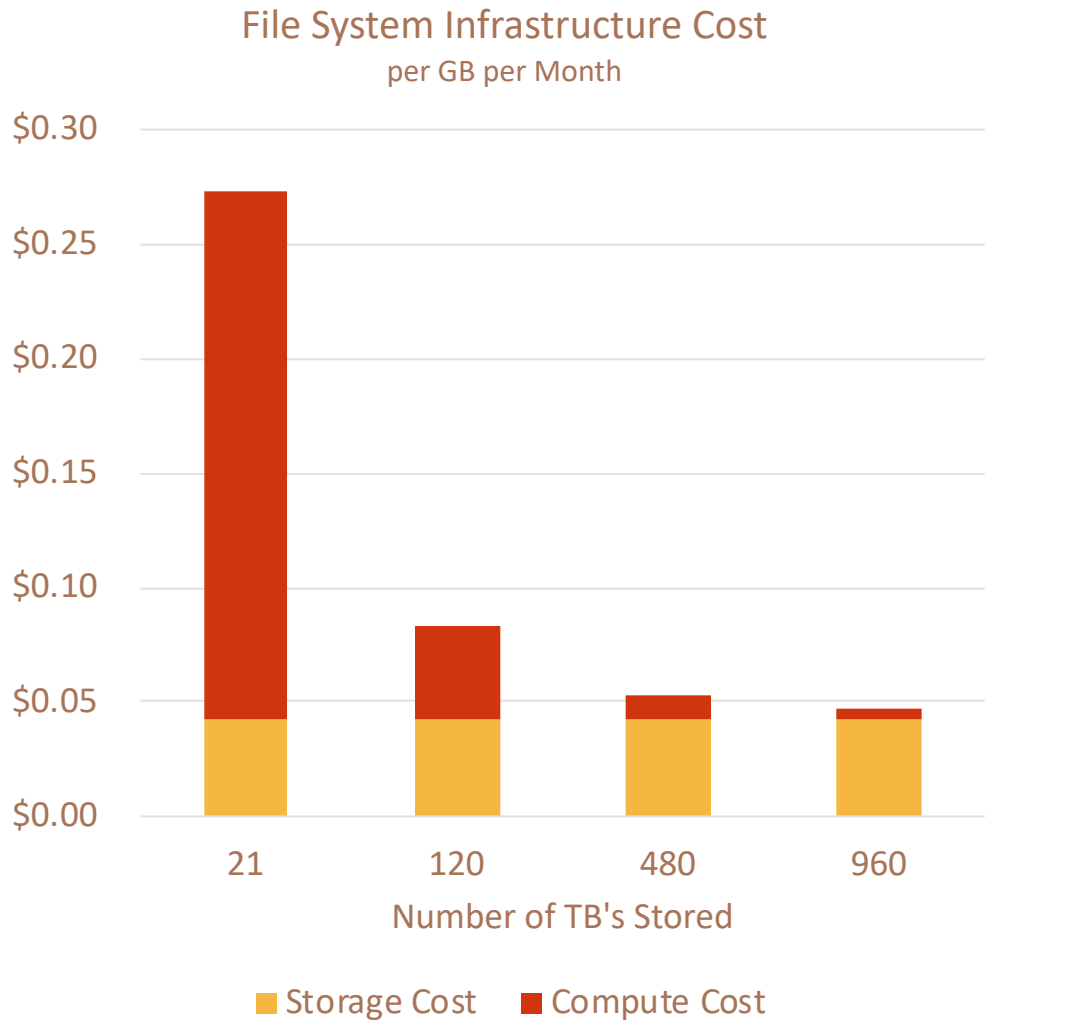

Performance – IO500 Benchmark

- Ranked 13th in the list of Fastest HPC filesystem
- List available at https://www.vi4io.org/io500

Competitive Cost Analysis – Why OCI ?

Cost/Performance Factor	OCI	AWS	GCP
Network Bandwidth / Compute	2 x 25 Gbps	1 x 25 Gbps	1 x 32 Gbps
Network Bandwidth for Block Storage	25 Gbps	14 Gbps	9.6 Gbps
Block Storage Cost (for 960TB filesystem)	$40,800	$236,720 (+480%)	$163,200 (+300%)
Bare Metal Compute Pricing (OCPU/Hour)	$0.0638	$0.126 (+97%)	$0.1184 (+85%)
Attach Block Volume to multiple Compute nodes	Yes	No	No
Egress charges (for 100TB)	$850	$8750 (+930%)	$8300 (+876%)

Storage Capacity

2 choices to scale up storage capacity

- Change the number of Block volumes attached to each ‘Building Block’
- Add more ‘Building Blocks’

Number of Building Blocks	Number of File Server Nodes	Min Storage Capacity (TB)*	Max Storage Capacity (TB)+
1	2	15	960
2	4	31	1,920
4	8	62	2,880
8	16	123	5,760
16	32	246	11,520

* For optimal throughput

+ For optimal throughput & Max storage capacity

Small Capacity File system also supported

Summary - IBM Spectrum Scale on Oracle Cloud Infrastructure

- Less than $0.05 GB/month infrastructure costs (Compute + Storage)
- IBM Spectrum Scale License - BYOL
- A HA scalable filesystem on OCI with POSIX semantics
- Tiering to OCI Object Storage
- Management GUI
- Dynamic storage capacity **independent** of storage throughput
- Mount to Windows or Linux
- **Automatically deploy** in minutes, no additional configuration required, https://github.com/oracle/oci-quickstart-ibm-spectrum-scale

Questions?